

ARTIFICIAL INTELLIGENCE AND THE FUTURE OF DEFENSE

STRATEGIC IMPLICATIONS FOR SMALL-
AND MEDIUM-SIZED FORCE PROVIDERS





HCSS helps governments, non-governmental organizations and the private sector to understand the fast-changing environment and seeks to anticipate the challenges of the future with practical policy solutions and advice.

This report is from the HCSS theme SECURITY. Our other themes are RESOURCES and GLOBAL TRENDS.

SECURITY

HCSS identifies and analyzes the developments that shape our security environment. We show the intricate and dynamic relations between political, military, economic, social, environmental, and technological drivers that shape policy space. Our strengths are a unique methodological base, deep domain knowledge and an extensive international network of partners.

HCSS assists in formulating and evaluating policy options on the basis of an integrated approach to security challenges and security solutions.



ARTIFICIAL INTELLIGENCE AND THE FUTURE OF DEFENSE:
STRATEGIC IMPLICATIONS FOR SMALL- AND MEDIUM-SIZED FORCE PROVIDERS

The Hague Centre for Strategic Studies (HCSS)

ISBN/EAN: 978-94-92102-54-6

Authors Stephan De Spiegeleire
Matthijs Maas
Tim Sweijts

This study is part of the HCSS *War Studies* Program, which gratefully acknowledges the financial support given by the Dutch Ministry of Defense. The views expressed here are solely those of the authors. .

© 2017 The Hague Centre for Strategic Studies. All rights reserved. No part of this report may be reproduced and/or published in any form by print, photo print, microfilm or any other means without prior written permission from HCSS. All images are subject to the licenses of their respective owners

Graphic design Ekaterina Novak

The Hague Centre for Strategic Studies

Lange Voorhout 16
2514 EE The Hague
The Netherlands

info@hcss.nl
HCSS.NL

ARTIFICIAL INTELLIGENCE AND THE FUTURE OF DEFENSE

STRATEGIC IMPLICATIONS FOR SMALL- AND MEDIUM-SIZED FORCE PROVIDERS

“One day the AIs are going to look back on us the same way we look at fossil skeletons on the plains of Africa. An upright ape living in dust with crude language and tools, all set for extinction.”

— Nathan Bateman in *Ex Machina*¹

“To bring about a revolution in military affairs, two things are normally needed: an objective development that will make it possible; and a man who will seize that development by the horns, ride it, and direct it.”

Martin van Creveld²

“Wie niet sterk is, moet slim zijn”

— Johan Cruijff³

1 Alex Garland, *Ex Machina*, Drama, Mystery, Sci-Fi, (2015).

2 Martin van Creveld, “Napoleon and the Dawn of Operational Warfare,” in *The Evolution of Operational Art: From Napoleon to the Present*, ed. John Andreas Olsen and Martin van Creveld (Oxford: Oxford University Press, 2011).

3 Louis van de Vuurst, “Johan Cruijff: Van Trainer Tot Adviseur” (Ajax.nl, Maart 2016), <http://www.ajax.nl/streams/ajax-actueel/johan-cruijff-van-trainer-tot-adviseur-.htm>

TABLES



Table of Contents

- Table of Figures** 8
- Table of Acronyms** 9
- Executive Summary** 12
- 1. Introduction** 20
- 2. What Is Artificial Intelligence?** 26
 - 2.1. Definitions.....26
 - 2.1.1. Intelligence 26
 - 2.1.2. Artificial Intelligence. 27
 - 2.1.3. Artificial Narrow-, General-, and Superintelligence 30
 - 2.2. A Brief History of AI..... 31
 - 2.2.1. Early Enthusiasm: The First AI Spring ('56-'75') 31
 - 2.2.2. 'Grandiose Objectives': The First AI Winter ('74-'80) 32
 - 2.2.3. Expert Systems: The Second AI Spring ('80-'87) 33
 - 2.2.4. 'Not the Next Wave'; The Second AI Winter ('87-'93) 33
 - 2.2.5. A Third, Sustained AI Spring ('93-'11) 34
 - 2.2.6. Big Data, Deep Learning, and an Artificial Intelligence Revolution 35
 - 2.3. AI: A Cookbook of Components, Approaches and Design Architectures..... 39
 - 2.3.1. Design Specifications: The Components of Intelligence 39
 - 2.3.2. Machine Learning: 5 Schools of Thought 40
- 3. AI – Today and Tomorrow** 44
 - 3.1. Core Technologies and Applications..... 45
 - 3.2. Markets amidst the Artificial Intelligence 'Insights Revolution'..... 47
 - 3.3. The Timing of AI Development: Future Scenarios..... 50
 - 3.4. Disruptors & Wild Cards..... 52
 - 3.5. Caveats: Legal and Ethical Concerns over AI..... 54
 - 3.6. Summing up – The AI Tipping Point..... 55
- 4. Defense – Yesterday, Today and Tomorrow** 61
 - 4.1. Epochs and Technologies..... 62
 - 4.2. Layers of 'Defense' 65
 - 4.2.1. Armed Force: AI and A2A 71
 - 4.2.2. Defense Organizations: AI and A2D (/D2D) 72
 - 4.2.3. Defense and Security Organizations: D2G (/G2G) 73

4.2.4. Defense and Security Ecosystem: E2I (/E2E)	74
4.3. AI in Other Armed Forces Today.....	77
4.3.1. China	77
4.3.2. Israel	80
4.3.3. Russia.....	81
4.3.4. US	83
4.4. Use Cases.....	87
4.4.1. Automating Cyber Operations (G2G, D2D)	87
4.4.2. Algorithmic Targeting (A2A)	88
4.4.3. Mission Handoff (A2A; A2D)	89
4.4.4. Situational Awareness and Understanding (A2A).....	90
4.4.5. Automated Planning and Manpower Allocation (A2A)	91
4.4.6. Target Systems Analysis/Target Audience Analysis (A2A, G2G, D2P).....	92
4.4.7. Lessons Learned – Operational and Non-Operational (A2A, A2D).....	93
4.4.8. Options Analysis for Capability Development (A2A, A2D)	93
4.4.9. Electronic Medical Records Analysis & Optimizing Medevac (A2A, D2G).....	94
4.4.10. Documents Classification and ‘Crypto-Preserving’ Intelligence Sharing	94
4.4.11. ‘P4’ Conflict Prevention – Predictive, Preventive, Personalized and Participatory ..	95
4.4.12. Empowering Societal Resilience (G2I, E2I).....	97
5. Conclusion.....	100
6. Bibliography	109
Appendix I: Asilomar AI Principles.....	137



Table of Figures

Figure 1: 4 layers of defense	13
Figure 2: Artificial Intelligence and 4 layers of defense.....	14
Figure 3: A Typology of different definitions of AI, and their underlying approaches.....	29
Figure 4: IBM's view of the history of computing	38
Figure 5: An overview of notable approaches and disciplines in AI and Machine Learning...	40
Figure 6: An overview of the competitive market landscape for machine intelligence.....	46
Figure 7: One scenario for an intelligence explosion.....	50
Figure 8: Possible AI development timelines	51
Figure 9: 4 layers of defense	68
Figure 10: Layers of interaction in the business world	69
Figure 11: Defense and security ecosystem.....	70
Figure 12: Shifts in the dominant models of thinking about purposive public action.....	72
Figure 13: Automated Planning and Manpower Allocation.....	92
Figure 14: Characteristics of the Conflict Model Vs the Resilience Model	98
Figure 15: Artificial Intelligence and 4 layers of defense copy	100



Table of Acronyms

A2A	Armed Force-to-Armed Force
A2D	Armed Force-to-Defense Organizations
AARs	After Action Report
ACTUV	Anti Submarine Warfare Continuous Trail Unmanned Vessel
AGI	Artificial General Intelligence
AI	Artificial Intelligence
ALPAC	Automatic Language Processing Advisory Committee
ANI	Artificial Narrow Intelligence
ARPA	Advanced Research Projects Agency
ASI	Artificial Superintelligence
ASW	Anti-Submarine Warfare
ATR	Automatic Target Recognition
AUVSI	Association for Unmanned Vehicle Systems International
B2B	Business-to-Business
B2C	Business-to-Customer
B2G	Business-to-Government
B2I	Business-to-Individual
BI	Business Intelligence
C2	Command and Control
CAGR	Compound Annual Growth Rate
D2C	Defense-to-Citizen
D2D	Defense Organization-to-Defense Organization
D2G	Defense Organization-to-Government
D2I	Defense Organization-to-Individual
D2P	Defense-to-Politics
D2S	Defense-to-Society
DARPA	Defense Advanced Research Projects Agency
DOD	Department of Defense
DSB	Defense Science Board
DSO	Defense and Security Organization
E2I	Ecosystem-to-Individual
EETN	Hellenic Artificial Intelligence Society
EMRA	Electronic Medical Record Analyzer
ERP	Enterprise Resource Planning
Fun LoL	Fundamental Limits of Learning



GCR	Global Catastrophic Risk
G2G	Government-to-Government
GDP	Gross Domestic Product
HPC	High-Performance Computing
HLMI	Human-level machine intelligence
ICT	Information and communications technology
ITS	Intelligent Tutoring Systems
IBM	International Business Machines Corporation
ISIS	Islamic State of Iraq and Syria
IDF	Israel Defense Forces
LAWS	Lethal Autonomous Weapon System
ML	Machine Learning
MIT	Massachusetts Institute of Technology
MCC	Microelectronics and Computer Technology Corporation
NLP	Natural language programming
OTS	Off-the-Shelf
PLA	People's Liberation Army
PC	Personal Computer
PT-AI	Philosophy and Theory of Artificial Intelligence
R&D	Research And Development
SMC	Small- to Medium-Sized Country
SME	Small- to Medium-Sized Enterprise
SAS	Statistical Analysis System
TAA	Target Audience Analysis
TNO	Nederlandse Organisatie voor Toegepast Natuurwetenschappelijk Onderzoek (Netherlands Organisation for Applied Scientific Research)
TRACE	Target Recognition and Adaptation in Contested Environments
TSA	Target Systems Analysis
USG	Universal Charging System
UAS	Unmanned Aerial System
UAV	Unmanned Aerial Vehicle
UGV	Unmanned Ground Vehicle
UUV	Unmanned Underwater Vehicle
X-risk	Existential Risk



EXECUTIVE SUMMARY



Executive Summary

Artificial intelligence (AI for short) is widely acknowledged to be one of the most dramatic technological game changers of our age. It is already starting – and is expected to continue – to have a disruptive impact on most walks of private and public life. Defense and security will be no exception. But how exactly will these be affected? Most current thinking on how AI will impact defense and security is incremental in nature. It assumes future armed forces that look very much like today's and continue to engage in essentially the same types of activities/operations with a very similar capability bundle. This is also what most current AI-related defense-specific R&D efforts as well as the first applications of AI in the military realm are focused on. This report argues that AI may have a more profound transformational impact on defense and security whereby new incarnations of 'armed force' start doing different things in novel ways.

Throughout history *intelligence* – the computational part of the ability to achieve goals in the world – has played a critical role in deciding who came out on top in various contests of wills. It is intelligence that allowed *homo sapiens* to fight its way to the top of this planet's food chain by crafting 'arms' and – especially – by developing intra-species cooperation strategies not just with family or kin members but also with strangers. It is intelligence that also allowed various groups of humans – first clans, then settlements, cities, (coalitions of) nation states, empires, etc. – to achieve a competitive advantage over others. Those groups that proved more intelligent and nimble in leveraging new physical and social technologies as they emerged typically ended up ahead of others.

We have now arrived at a juncture in human development where *homo sapiens* is starting to uncover pathways towards developing novel forms of 'artificial' intelligence that may prove equal or even superior to our own. Our very own evolutionary history suggests that this may once again upset the balance of power. Turbulent advances in AI – especially in the past 2 years with the ascendance of various deep learning tools – are already uprooting many traditional practices across different domains of human activity. Examples include scientific discovery, machine translation, digital assistants, search, mobility, health, etc.

Against the background of these stormy developments this study examines the implications of AI for defense and security organizations (DSOs). It juxtaposes 3 different layers (/generations) of AI against 4 different layers (and possibly generations) of Armed Force. For AI, it differentiates between **Artificial Narrow Intelligence** (ANI) – machine intelligence that *equals or exceeds* human intelligence for *specific* tasks; **Artificial General Intelligence** (AGI) – machine intelligence *meeting the full range* of human performance across *any* task; and **Artificial Superintelligence** (ASI): machine intelligence that *exceeds*

human intelligence across *any* task. ANI is already amongst us in various guises – often unnoticed, as in our search engines or the assistants on our mobile phones. The timing of the advent of AGI remains controversial, but most experts expect it by the middle of this century – within the time horizon of current long-term defense planning. Most also expect ASI to emerge relatively quickly thereafter, although few anticipate this to spark real artificial consciousness (yet).

Against these three layers of AI this report plots 4 layers of defense. These can be thought of as concentric circles around a **first (core) layer**: our **armed forces** themselves. These are the active agents in terms of wetware (humans), hardware and software behind our efforts to generate defense and security effects. The **second layer** consists of the **supporting defense organization** we have built around that ‘sharp edge’ of operators in order to enable them to function properly – the various military support entities and the Ministry of Defense. This report emphasizes that the current way in which most nations have organized this layer really only emerged in the industrial age and may become subject to the same disruptive evolutionary pressures that spawned it in the first place.




FIGURE 1: 4 LAYERS OF DEFENSE

We have already seen this in the **third layer**, which positions defense as a player in an increasingly more **whole-of-government** security-oriented approach in which all the entities that are responsible for defense and security across our governments start behaving and acting more coherently: in which also they are the main actors. The **fourth** and **final layer** is one we label the **defense and security ecosystem** in which the first, second and third layers become (potential) catalysts or even a platform for a much broader defense and security ecosystem of sensors and effectors that can be actuated in order to achieve defense and security goals. Mapping these 4 layers of ‘defense’ against the 3 layers of AI yields the following table, which synthesizes the options available to small- and medium-sized force providers.

	Artificial Narrow Intelligence	Artificial General Intelligence	Artificial Super Intelligence
'Armed Force'	Monitor frontrunners and acquire opportunistically	Review robustness current force structure	
Ministry of Defense	Identify short-term challenges and opportunities	Identify long-term challenges and opportunities	
Comprehensive	Identify short-term challenges and opportunities	Identify long-term challenges and opportunities	Catalyze dialogue with all stakeholders
Defense and Security Ecosystem	Explore new niches	Explore new niches	Existential challenges and opportunities: fundamental rethink of defense

FIGURE 2: ARTIFICIAL INTELLIGENCE AND 4 LAYERS OF DEFENSE

Most of the current literature on AI and defense focuses on the **top-left cell of this table: artificial narrow intelligence** that helps current operators and warfighters accomplish (essentially) their current tasks within the current (industrial-age) paradigm. This study surveys how some of the key military players today – none more than the United States, but other aspirational peer-competitors such as China and Russia as well – are focusing on this cell in their defense AI efforts. We argue that – based on the recent record of cost inflation in this cell – the financial implications of focusing primarily on this particular type of defense AI may very well prove to be exorbitant. Especially US (and/or UK) investments in defense AI, competing with Chinese (and possibly Russian) efforts, can be expected to be both substantial and significant. This means that they will yield uniquely new and powerful – in the industrial-kinetic sense – capability options. Even if these new capabilities are expensive, they may still prove to provide better value-for-money than anything other European force providers (certainly from small- to medium-sized countries - SMCs) might be able to generate in this space in their own right. The study therefore recommends a cautiously pragmatic and opportunistic attitude towards this top-left cell of the AI option-space. Investment opportunities will undoubtedly emerge in this segment that small- to medium-sized defense providers might be able to jump on – even based on their own R&D funds; but our suggestion here is to reap the potential benefits of the R&D investments of the frontrunners, as our defense industrial base has done in many recent 'big ticket' defense procurement projects.




As the overall table, and especially the right-downwards orange arrow in [Figure 2](#), suggests, however, we anticipate significantly superior value-for-money opportunities for smaller DSOs in the other cells of this table – certainly downwards towards the defense and security ecosystem, but arguably also towards the right bottom as we move towards **artificial super intelligence**.

In the first instance, we suggest that smaller force providers point their attention and **focus downwards** in this table – towards the ecosystem layer of defense. In the first layer, issues like classification/secretcy and the *de facto* oligopolistic market structure lead to highly politicized – and thus expensive – military capability and acquisition choices. In the **second layer** the dynamics of the market for defense solutions seem much less atypical than in the first layer and already much more dual-use. We therefore suspect that many functions currently being fulfilled by civil servants in ministries of defense (HR, procurement, inventory management, but increasingly also policy analysis and making etc.) may increasingly be encroached upon by various (learning) AI algorithms, which will spill over from the civilian sector. As in the private sector, the job displacement implications of this trend will have to be carefully managed; but we think that the likely efficiency gains will make this transition irresistible and worthwhile.

This is even more the case in the **third layer** of [Figure 2](#) – the whole-of-government one – where we expect an enormous market for learning algorithms to improve (and rationalize) the management and coherence of the massive amount of paperwork our bureaucracies generate. The shift will be to dramatically cut down on the various transaction and coordination costs our governments currently incur (inside-out), and to start building, curating and exploiting a knowledge base built on more pragmatic learning from the security effects that are generated by the various security-relevant actions across government (outside-in).


Some of the most promising defense and security AI applications are in our opinion likely to emerge in and be focused on the **fourth layer** of defense: the defense and security ecosystem layer. We describe in the report how in the private sector, the world's largest companies are in the process of morphing themselves from pipeline businesses into platform businesses at the heart of much broader ecosystems that co-create value for both the companies and the ecosystem partners. AI is playing an increasingly powerful role in this transformation through a combination of deep learning, large datasets and high performance computing. We anticipate a similar trend in defense and security. Some important players in this field, such as China and Russia, are already embracing and exploiting opportunities in this layer to further what they perceive to be their defense and security interests. It still continues to be the case, however, that most of the world's most cutting-edge AI research and start-ups are located in NATO nations. Even China's Baidu has a strong presence in Silicon Valley with some leading AI researchers, led (until recently) by London-born Chinese-American AI luminary Andrew Ng. A more systematic dialogue between our DSOs and the AI community may therefore prove to be a critical ingredient in providing unique defense value for defense money, especially – we suspect – in the areas of prevention and resilience.



Alongside moving our focus downwards in [Figure 2](#), this study also recommends that defense and security organizations **move their efforts towards the right of that figure**. As AI moves towards more **artificial general intelligence**-types of comprehensive solutions, the still relatively neat lines between defense and non-defense technologies and applications that we still see today are likely to blur even more. This will work both ways. On the one hand, non-defense and security AI-augmented technologies and applications will increasingly have to worry about security and privacy concerns and will have to evolve into more sustainably reliable solutions. This might also prove useful for defense and security organizations (and – maybe increasingly – defense and security ecosystems). Quantum- and/or blockchain-based solutions, once they come online, may play an important role in this. On the other hand, more purely defense and security AI-augmented technologies and applications will increasingly be able to piggyback on, and perhaps even merge, with the core technologies and applications developed by their civilian counterparts for totally unrelated purposes. To give a concrete example: as medicine starts moving more and more towards what some now call P4 medicine (personalized, predictive, preventative, participative), the very sensors, data and algorithms that will be used for health purposes might prove equally useful for defense and security purposes. Obtaining societal approval for a more proactive role in this area by more platform-like governments with a broader ecosystem around them, will prove of critical importance. This will require that the privacy, legal and ethical implications be adequately addressed in open(-minded) discussions between all of the stakeholders.

Even if we may still be some time away from developing ‘full’ or ‘general’ AI, it is important that we already start thinking about the legal and ethical implications, and considering measures for the responsible supervision, regulation, and governance of the design and deployment of AI systems. We need to heed the call by prominent AI researchers to commit to the “careful monitoring and deliberation about the implications of AI advances for defense and warfare, including potentially destabilizing developments and deployments”. The study suggests that the new Asilomar principles from the Conference on Beneficial AI (Appendix 1) provide an excellent starting point for such a discussion. It also suggests that the increased use of AI in the prevention and resilience realms may lend itself to more constructive and participatory discussions than the debates so far on autonomous weapon systems or ‘killer-drones’.

The final – existentially important – cell in this option space is the one on the bottom right. Artificial super-intelligence – i.e. intelligence that is superior to that of *homo sapiens* – is likely to pose quite unique challenges to defense and security planners. Our *longue durée* historical overview of the nexus between intelligence and defense has shown how the non-physical, purely cognitive part of human force proved evolutionarily superior to the sometimes physically stronger but cognitively weaker force of other species. If that history is a lesson, then the emergence of super-human intelligence is surely a transition that requires extraordinarily careful human consideration. We agree with many other analysts who have looked into this issue that this is an unusually – arguably unprecedentedly – wicked challenge.




Both artificial narrow intelligence and artificial general intelligence are likely to open up many unique opportunities that may enhance our human – individual and collective – experience. Many – maybe even most – of us are likely to welcome that part of it in our lives. At the same time, however, we also concur with the majority of analysts who anticipate major – maybe even existential – threats when non-human intelligence exceeds its human equivalent and turns into artificial superintelligence. Who will be *homo sapiens*' custodian or ambassador in this debate? Can we imagine a situation where governments worldwide team up together behind the flag of *homo sapiens* in order to address the AI control issue, at the same time as they are gearing up their own AI capabilities? Small- and medium- sized powers may actually be well positioned to take the initiative for such an effort.

Our own summary assessment, therefore, is that developed small- to medium-sized force providers are better served with a capability, partnership, concept, and policy portfolio that is focused more towards the lower and the right (and – depending on the observable future trajectory of the trend towards artificial super intelligence – bottom-right) regions of this options space, than towards the industrial-kinetic top-left one. They may still be able to benefit from the investments their bigger allies are likely to make in that part of the option space. At the same time, however, we suggest that especially for small- to medium-sized DSOs the other regions of this option space present superior value for money from a defense and security point of view.

Pivoting towards those regions requires fundamental changes in how DSO see themselves, in what they do and how they do it. Traditional approaches like the much-vaunted 'golden triangle', or the focus on enhancing industrial-kinetic capabilities will no longer suffice. The most important - and difficult - shift, we suspect, lies in redefining and redesigning 'defense' itself. Only DSOs that will be able to secure advantageous positions within the quickly and dynamically expanding AI ecosystem will be in a position to make pro-active instead of re-active choices in this field. The report identifies great opportunities for such truly pro-active choices in the areas of prevention and resilience - both at home and abroad. through which AI can start providing the types of sustainable security and defense objectives that our populations and political leaders expect in different ways.

This brings us to the following set of recommendations for small- to medium-sized defense and security organizations:

1. 'Open the windows' on how we think about future applications of AI to different (also possibly disruptive) futures of defense and security – and do so not from the point of view of our current mindset, structures, capabilities, partnership choices, etc.; but from the point of view of the broader defense and security objectives we want to sustainably achieve;
2. Identify a prioritized list of AI applications that can be rapidly implemented in the framework of existing structures and processes ('quick wins'). This effort would benefit from an agent of change that can start initiating, experimenting with and stimulating the implementation of (some of) these applications within



the defense organization – especially those that focus more on prevention and resilience capabilities.

3. Start a dialogue with various entities across the entire defense and security ecosystem about both promising opportunities and critical challenges in AI (including legal, ethical and privacy aspects) throughout the option space;
4. Start designing a capability, partnership, concept and policy portfolio that is focused on the lower regions of this defense AI options space;
5. Get and remain situationally aware of what is going on in this exceptionally dynamic field. Acknowledge that this cannot be done the 'old' triangular (golden or not) way, but requires engaging the entire defense and security ecosystem;
6. Make sure that new emerging and promising options which may not (yet) have advocates within our defense and security organizations stand a chance against the bigger-ticket legacy items for which there are (powerful) stakeholders;
7. Start thinking about artificial superintelligence and engage with the community that has started thinking about actionable options in that part of the option space (also recognizing that this may open up new avenues for engaging with actors such as China and/or the Russian Federation).

We are encouraged by the fact that the Dutch DSO has already started moving in this direction. The importance of both artificial intelligence and ecosystems has been highlighted in the department's Strategic Knowledge and Innovation Agenda (SKIA) 2016-2020 and in its multi-year perspective on the future of the armed forces (*Houvast in een onzekere wereld*). The concept of the 'defense and security ecosystem' is increasingly being embraced and internalized by the organization, as most vividly illustrated by the 2017 Future Force Conference that was co-organized by the Minister of Defense and the Chief of Defense and was entirely devoted to the operationalization of this concept. Building on these first steps, we sincerely hope this report contributes to a more fundamental rethink of how AI may affect the future of defense and security and how our DSOs can responsibly harness developments in this field to achieve more sustainable defense and security solutions.

1 INTRODUCTION

1. Introduction

Nature has developed multiple strategies that allow biological organisms to *defend* themselves against other members of their own species, against other species and against nature. In the inter-species offense/defense battle humans – homo *sapiens* – appear to be among the most vulnerable of all major contestants. Our size is much smaller than that of the few remaining large animals like bears or elephants. Our sheer physical strength is dwarfed by that of most other primates or other much stronger species. We have little in terms of natural force protection – no armature, no spikes, etc. We are not naturally exceptionally mobile – just try outrunning a cheetah or even a small dog. Our bodies have not developed special weapons like venom, sharp teeth or claws, vicious horns, electric rays or any of the other types that nature has invented.⁴ Our young are virtually helpless against all sorts of potential predators without parental assistance for a much longer time than any other known species.⁵

And yet since the *Cognitive Revolution* (~70,000 years ago), the ‘armed force’, the capability bundle that humans were able to evolutionarily develop, ended up overpowering any and all of these better protected, faster, bigger, stronger, intrinsically deadlier species. It was not brute physical force that bestowed homo *sapiens* this decisive evolutionary advantage. It was the intelligence – in the broader sense of that term – that at some critical point in our evolution was sparked in the human brain’s neural networks, that propelled homo *sapiens* from an ‘animal of no significance’ onto the undisputed top of the food chain.⁶

Human intelligence achieved this impressive feat in two ways. First of all, humans used their cognitive skills to craft lethal tools that compensated for their intrinsic physical frailty. The human species is not alone in manufacturing and using tools,⁷ but the extent

4 For some telling examples, see Jessica Hullinger, “7 Absolutely Insane Animal Defense Mechanisms,” *Mental Floss*, October 6, 2013, <http://mentalfloss.com/article/12258/7-absolutely-insane-animal-defense-mechanisms>.

5 And the specificities of our brain – the source of our intelligence – may be a factor in this, see Holly M. Dunsworth et al., “Metabolic Hypothesis for Human Altriciality,” *Proceedings of the National Academy of Sciences* 109, no. 38 (September 18, 2012): 15212–16, doi:10.1073/pnas.1205282109. Kate Wong, “Why Humans Give Birth to Helpless Babies,” *Scientific American Blog Network*, August 28, 2012, <https://blogs.scientificamerican.com/observations/why-humans-give-birth-to-helpless-babies/>; Karen R. Rosenberg and Wenda R. Trevathan, “Are Human Infants Altricial?,” 2015, <http://meeting.physanth.org/program/2015/session15/rosenberg-2015-are-human-infants-altricial.html>.

6 That point in time needs to be taken figuratively as intelligence has evolved incrementally over time. Cf. Michael Harré, “Social Network Size Linked to Brain Size,” *Scientific American*, August 7, 2012, <http://www.scientificamerican.com/article/social-network-size-linked-brain-size/>; Yuval N. Harari, *Sapiens: A Brief History of Humankind*, First U.S. edition (New York: Harper, 2015); Ian Morris, *War! What Is It Good for? Conflict and the Progress of Civilization from Primates to Robots* (New York: Picador USA, 2015)

7 For some fascinating examples, see Robert W. Shumaker, Kristina R. Walkup, and Benjamin B. Beck, *Animal Tool Behavior: The Use and Manufacture of Tools by Animals* (JHU Press, 2011). Of particular note is that apparently only humans and chimpanzees have the foresight and intellectual ingenuity to make tools to stab at other animals. (p. 217). See also Jill D. Pruetz and Paco Bertolani, “Savanna Chimpanzees, Pan Troglodytes Verus, Hunt with Tools,” *Current*

to- and ingenuity with which it did this does set it apart from all other species. But there is second way in which human intelligence thrust us into our current position of evolutionary dominance. Using the very same neural networks in our brain that crafted more powerful implements of brute physical force, this second application of intelligence worked its unique magic through more subtle, but ultimately far more powerful pathways. In his impressive survey of the history of *Homo Sapiens*, the Israeli (military) historian Yuval Noah Harari reminds us that “at the time of the Cognitive Revolution, the planet was home to about 200 genera of large terrestrial mammals weighing over 100 pounds. At the time of the Agricultural Revolution, only about a hundred remained. Homo sapiens drove about half of the planet’s big beasts to extinction, long before humans invented the wheel, writing, or iron tools.”⁸ Homo sapiens proved equally ruthlessly successful with other advanced great apes such as *homo soloensis*, *homoe denisova*, *homo florensis* and the more widely known *homo neanderthalensis*, who became extinct once (or not long after) they shared the stage with *homo sapiens*.

There is an increasing amount of evidence that the extinction of all of these other mammals was not due to their inferior physical strength or lack of brains (*homo neanderthalensis*, in fact, was both taller, stronger *and* had bigger brains than *homo sapiens*). Even the tools that homo sapiens manufactured to give itself at least a fighting chance against its opponents did not make the crucial difference. The decisive factor seems to have been the unique ability that *homo sapiens* possessed to leverage his (and her) cognitive abilities⁹ to form social relationships and share accumulated knowledge. It is through culture that it proved able to develop flexible intra-species cooperation arrangements, not just with family or kin members but also with strangers.¹⁰ The underlying form of ‘armed force’ that enabled humans’ evolutionary ascendance consisted much less of brute force physical technologies than of more subtle and cunning individual and – especially – social cognitive technologies. As Harari observes “we are successful because we can form long-lasting relationships with many others in diverse and flexible ways, and ... this, combined with our native intelligence, explains why homo sapiens came to dominate the planet.”¹¹

When we turn our attention from inter-species to intraspecies defense mechanisms and practices – i.e. between different clans of homo sapiens – we encounter a strikingly


Biology 17, no. 5 (March 2007): 412–17, doi:10.1016/j.cub.2006.12.042. - who point out that chimpanzees who have been observed hunting with tools include females and immature chimpanzees, suggesting the same may have been the case with early hominids.

8 On the Holocene (or Anthropocene or Sixth) extinction, see Harari, *Sapiens*, 2015. Elizabeth Kolbert, *The Sixth Extinction: An Unnatural History*, First edition (New York: Henry Holt and Company, 2014); Ashley Dawson, *Extinction: A Radical History*. (New York: Or Books Llc, 2016).

9 Including the ability to develop and draw on an innate Theory of Mind, which is the ability to attribute mental states (e.g. beliefs and desires) to other people in order to predict their behavior. Cf. Robert M. Seyfarth and Dorothy L. Cheney, “Affiliation, Empathy, and the Origins of Theory of Mind,” *PNAS* 110, no. 2 (2013), http://m.pnas.org/content/110/Supplement_2/10349.full.pdf; Marie Devaine, Guillaume Hollard, and Jean Daunizeau, “Theory of Mind: Did Evolution Fool Us?,” *PLOS ONE* 9, no. 2 (February 5, 2014): e87619, doi:10.1371/journal.pone.0087619.

10 Cf Yuval N. Harari, “Why Humans Run the World,” *Ideas.ted.com*, June 16, 2015, <http://ideas.ted.com/why-humans-run-the-world/>.

11 Harré, “Social Network Size Linked to Brain Size” *Scientific American*; Joanne L. Powell et al., “Orbital Prefrontal Cortex Volume Correlates with Social Cognitive Competence,” *Neuropsychologia* 48, no. 12 (October 2010): 3554–62, doi:10.1016/j.neuropsychologia.2010.08.004. Similar evolutionary advantages also accrued to other marginally more intelligent (and not only ‘stronger’) species at lower ranks of the food chain.



similar duality. On the one hand, human clans applied their superior cognitive abilities – with increasing vigor and creativity – since at least the Mesolithic period (~8,000 BC) to competitively forge ever more destructive physical capacity bundles.¹² These peaked in the late Industrial Age with the development of nuclear weapons.¹³ Groups of humans used this ingenuity to prevail in direct physical confrontation with the militaries of other groups of humans. At the same time, however, rather than just creating lethal weaponry, humans also leveraged their intelligence in various non-physical ways to become ever more sophisticated in how they prepare for war, including, but not limited to, such fields as propaganda, intelligence, communications, command and control, and logistics. In current-day parlance, human armed force has always been *hybrid* in this sense.

We have now arrived at a juncture in evolution where this very same *homo sapiens* is starting to uncover pathways towards developing novel forms of intelligence that may prove equal or even superior to its own. The recent irruption of various forms of artificial intelligence – AI for short – is already uprooting many traditional practices across different domains of human activity. These include, to name but a few examples, the automotive industry (from stock and inventory management to driverless cars), health sciences (from sequencing care interventions to advanced melanoma screening and detection), financial trading (algorithmic asset management), and advertising (behavioral targeting). These examples are merely a few applications of AI that are visible today. Based on historical experience with the development of new technologies, the emergence of AI will not just affect our security environment, but also the way in which our defense and security organizations (DSOs) operate.¹⁴ It will in one way or another shape the roles, responsibilities and future capabilities of choice of DSOs.

In this context, it is neither a secret nor a surprise that DSOs around the world are scrambling to pre-position themselves in the emerging field of AI in order to stay ahead, leapfrog or simply keep up with potential developments. The US is heading the pack spending up to \$18 billion USD in the context of its much acclaimed Third Offset strategy,¹⁵ a significant part of which is dedicated to furthering the development of AI applications, but other states, including China, Russia and Israel, are following suit.

At the same time, it is important to consider legal and ethical principles and guidelines. Like all new technologies, it is inevitable that AI will be deployed and leveraged, one way or another, by many different actors, for a wide range of goals and purposes – not all of which are compatible with or beneficial to mankind. Moreover, the various uses and impacts of AI will not occur independently of one another, or of a multitude of other

12 Azar Gat, *War in Human Civilization* (Oxford University Press, 2008), <http://www.oupcanada.com/catalog/9780199236633.html>, 26-31. If not earlier: the bow, for instance, was invented ~20,000 years ago, but may have been merely used for hunting.

13 David E Hoffman, *The Dead Hand: The Untold Story of the Cold War Arms Race and Its Dangerous Legacy* (New York: Anchor Books, 2009), <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=7359>

14 For an excellent overview of the interaction between social and technological developments and the organized 'pursuit of power', see William H. McNeill, *The Pursuit of Power: Technology, Armed Force, and Society since A.D. 1000*, 1 edition (Chicago: University of Chicago Press, 1984).

15 Andrea Shalal, "Pentagon Eyes \$12-15 Billion for Early Work on New Technologies," *Reuters*, December 14, 2015, <http://www.reuters.com/article/us-usa-military-technologies-idUSKBN0TX1UR20151214>

societal and technological developments.¹⁶ Already today we are seeing the societal fallout which results from the so-called 'algorithmic bias',¹⁷ and in 2016 we have all seen first-hand the consequences of weaponized 'computational propaganda'¹⁸ on democratic processes and public trust.

Against the background of these stormy developments, this study considers the implications of AI for DSOs. What role will, can and should DSOs play in leveraging the opportunities and defending against the threats? What are potential effects for who they are, what they do, and how they do it? What are the legal and ethical challenges associated with the emergence and deployment of AI – and how do we address these? And what is the appropriate course of action for smaller defense providers, in particular the Netherlands: should they simply piggyback on the efforts of other larger defense organizations that possess greater financial means to make a difference in the AI domain (that is developing much more dynamically outside of defense R&D than inside of it)? Or are there still some niches where they can provide unique added value? These are the questions this paper seeks to address.

The authors of this think-piece are not technical AI specialists. The focus is therefore not primarily on the various technologies that undergird the extraordinarily dynamically developing field of AI – although we were fortunate to be able to rely on a number of excellent overview pieces in this area. We are also not military professionals, defense planners, or acquisition practitioners, which puts us both at a relative disadvantage, but also – we would submit – at a relative advantage vis-a-vis those groups, because we are free(-er) to think outside of the confines of current portfolios and prevailing notions of what armed force is and should be. What we do try to bring to the table, however, are a number of broader strategic considerations of what the upcoming AI revolution will mean for what defense is all about.


The body of publicly available literature about AI in the military-technological realm is relatively small and – from our point of view – remarkably conventional. Contrary to a number of other disciplines (e.g., health, education or finance), where the general

16 Stanford University, "Artificial Intelligence and Life in 2030 One Hundred Year Study on Artificial Intelligence I Report of the 2015 Study Pane," *Stanford One Hundred Year Study on Artificial Intelligence (AI100)*, September 2016, https://ai100.stanford.edu/sites/default/files/ai_100_report_0901fnlc_single.pdf.

17 Lauren Kirchner, "When Discrimination Is Baked Into Algorithms," *The Atlantic*, September 6, 2015, <https://www.theatlantic.com/business/archive/2015/09/discrimination-algorithms-disparate-impact/403969/>.

18 Cf. Craig Timberg, "Russian Propaganda Effort Helped Spread 'fake News' during Election, Experts Say," *Washington Post*, 2016, https://www.washingtonpost.com/business/economy/russian-propaganda-effort-helped-spread-fake-news-during-election-experts-say/2016/11/24/793903b6-8a40-4ca9-b712-716af66098fe_story.html.

For academic analysis of computational propaganda, see Samuel C. Woolley and Philip N. Howard, "Automation, Algorithms, and Politics| Political Communication, Computational Propaganda, and Autonomous Agents – Introduction," *International Journal of Communication* 10, no. 0 (October 12, 2016): 9., Oxford Internet Institute, *Video: Politics, Propaganda, and Bots–The Changing Nature of Cyber Warfare*, 2016, <https://www.oii.ox.ac.uk/video-politics-propaganda-and-bots-the-changing-nature-of-cyber-warfare/>., Samantha Shorey and Philip N. Howard, "Automation, Algorithms, and Politics| Automation, Big Data and Politics: A Research Review," *International Journal of Communication* 10, no. 0 (October 12, 2016): 24.. For a journalistic angle, see Berit Anderson and Brett Horvath, "The Rise of the Weaponized AI Propaganda Machine," *Scout*, 2017, <https://scout.ai/story/the-rise-of-the-weaponized-ai-propaganda-machine>. and Tony Aubé, "AI and the End of Truth," *The Startup*, February 14, 2017, <https://medium.com/swlh/ai-and-the-end-of-truth-9a42675de18#drxlxqv4z>.



expectation is that AI will bring fundamental and even disruptive changes to the essence of what these areas of human activity are all about, the military focus on AI today remains primarily anchored within the current operating paradigm of fighting and defeating enemies that range on some continuum between ‘conventional’ and ‘guerilla’¹⁹. This paradigm precludes the ability to consider alternative yet plausible and promising future incarnations of armed force and armed forces. We therefore decided that the most useful contribution we could make in this area was to try and open the windows and let in some fresh air by taking a broader and – we hope – more fundamental, effects-based, future-oriented and strategic look at what AI might mean for DSOs.

The study is structured as follows. We start chapter two, *What Is Artificial Intelligence?*, by considering definitions of ‘intelligence’ and of ‘artificial intelligence’; the history of the technology, and its main components, disciplinary fields and schools of design. In chapter three, *AI – Today and Tomorrow*, we proceed to take stock of the state of the art in the discipline of AI: where the field stands today; what its currently most promising applications are; what the AI market looks like and what we can – and cannot – sensibly say about the future of AI. We conclude that chapter with some thoughts about what we see as three different generations of emergent AI. These initial chapters two and three are written as a primer on AI – and readers who feel familiar with these issues can skip ahead.

Chapter four, *Defense – Yesterday, Today and Tomorrow*, presents a similarly evolutionary analysis of defense and armed force, and proposes a new conceptual framework to consider the impact of AI on defense. This framework is used to explore how different generations of AI might feed into different generations of defense which may promise not just more cost-effective defense solutions – offering better defense and security value for defense and security money – but also enable more accurate and preventive forms of force projection. This is because they offer greater opportunity to utilize non-lethal force and avoid collateral damage, thereby enabling a way of war that could, perhaps paradoxically, prove more humane, not less. The *Conclusion* then offers a synthesis, considers the role of smaller force providers such as the Netherlands, and offers a set of recommendations for how they can best prepare for the advent of AI.

19 See Jeffrey A. Friedman, *The 2006 Lebanon Campaign and the Future of Warfare: Implications for Army and Defense Policy* (Lulu.com, 2011). These terms are conditional - and can be (and are - in the literature) substituted by many others. Our main point here is mainly that military strategists tend to focus on countering the ‘military’ effort of [different categorizations of the possible or real] enemy, as opposed to on achieving our strategic defense and security goals. An analogy we have found useful to convey this point is fire-fighting. Outside of war, modern cities in the developed world do not burn any more - as they used to do regularly and devastatingly in the past, even up to the early 20th century. This is partly due to more efficient firefighting tools (the idea that cities should invest in these emerged in wealthy 17th century Amsterdam, see Stephen J. Pyne et al., *Flammable Cities: Urban Conflagration and the Making of the Modern World* (University of Wisconsin Press, 2012). But it has more to do with intelligent prevention techniques like building codes, regulations and the like.



2 WHAT IS ARTIFICIAL INTELLIGENCE?

- 2.1 **Definitions**
- 2.2 **A Brief History of AI**
- 2.3 **AI: a Cookbook of Components, Approaches and Design Architectures**

2. What Is Artificial Intelligence?

2.1. Definitions

“By far the greatest danger of Artificial Intelligence,” AI theorist Eliezer Yudkowsky has observed, “is that people conclude too early that they understand it.”²⁰ As such, it is important to first properly delineate what AI is – or can be. To do so, we must first grasp what *general* intelligence refers to.

2.1.1. Intelligence

Intelligence is a cornerstone of the human condition – we even named ourselves after it.²¹ Intelligence is what allows you to understand (derive meaning from) the sentences you are reading here. It is a mental activity that we, as human beings, are uniquely qualified to exercise. Different elements of and processes within your brain, making use of other parts of your biology (e.g. your eyes), work together to make possible the act of reading and comprehension.

Our understanding of the etymological roots of the word intelligence is somewhat murky, but we do know it is closely related to the activity you are now engaging in – to read and understand. The word consists of two parts: the prefix *inter*, meaning ‘between’ and the Latin verb *legere*, meaning initially ‘to choose, to pick out’ and from there evolving into the meaning ‘to read’.²² Intelligence thus implies to gather, to collect, to assemble or to choose, and to form an impression, thus leading one to finally understand, perceive, or know.

Already in the 15th century, the word intelligence was understood as “superior understanding, sagacity, quality of being intelligent”.²³ The sense of “information received

20 Eliezer Yudkowsky, “Artificial Intelligence as a Positive and Negative Factor in Global Risk,” *Global Catastrophic Risks* 1 (2008): 303., pg 1. The article also appears in the excellent compilation Nick Bostrom and Milan M. Cirkovic, *Global Catastrophic Risks*, 1 edition (Oxford; New York: Oxford University Press, 2011)

21 The term ‘homo sapiens’ means ‘wise man’. For a wonderfully witty tongue-in-cheek comment by Google’s Director of Engineering, Open Source and Making Science on humans’ narcissistic self-perception, see Chris DiBona, “The Limits of Biological Intelligence,” in *What to Think about Machines That Think: Today’s Leading Thinkers on the Age of Machine Intelligence*, ed. John Brockman (HarperCollins, 2015).

22 Douglas Harper, “Online Etymology Dictionary: Arm,” accessed September 20, 2016, www.etymonline.com/index.php?term=arm. The Latin word in turn evolved from a proto-Indo-European root **leg-* (1), meaning “to pick together, gather, collect”. Reading, therefore, has been etymologically conjectured to imply ‘to pick out words’. This PIE root **leg-* (1) also makes it closely related to the Latin word *logos* “word, speech, thought”, which we know from the various scientific disciplines like sociology, anthropology, microbiology, etc.

23 Douglas Harper, “Online Etymology Dictionary: Intelligence,” accessed September 28, 2016, http://www.etymonline.com/index.php?allowed_in_frame=0&search=intelligence.

or imparted, news” was first recorded in the mid-15th century, from which time also stems its other frequent meaning – also in diplomatic and military circles – of “secret information from spies” (1580s).²⁴ Today, intelligence means different things to different people, both in daily parlance and across various academic disciplines.²⁵

Usage ranges from fairly quotidian comparisons of inter-human skill or markers of authority (focusing on the social function of intelligence), to extremely theoretical definitions of intelligence as just one instantiation of a universal class of “optimization processes” which, in the broadest sense, includes even the adaptive mechanisms of natural selection.²⁶ A comprehensive survey of over 70 definitions of ‘intelligence’, conducted by Legg & Hutter, argues that it is hard to settle on a single ‘correct’ term, however they do note that many of the most concise and precise definitions share a number of features: (1) intelligence is a property of some agent that interacts with an environment; (2) intelligence is generally indicative of that agent’s ability to succeed at a particular task or stated goal; (3) there is an emphasis on learning, adaptation, and flexibility within a wide range of environments and scenarios.²⁷

Combining these features, they distill a fundamental definition of intelligence, as denoting:

“an agent’s ability to achieve goals in a wide range of environments.”²⁸

For the purpose of this paper, we will operationalize this into the more specific and functional definition coined by Stanford University’s Formal Reasoning Group, which usefully covers both ‘natural’ (human and animal) and ‘artificial’ forms of intelligence:

“Intelligence is the computational part of the ability to achieve goals in the world.”²⁹

From this definition of intelligence, which refers to both internal processes (“computation...”) that act in the service of bringing about external results (“...the ability to achieve goals...”) across complex, dynamic environments (“...in the world”), we can proceed to home in on a definition of artificial intelligence.

2.1.2. Artificial Intelligence

In its broadest sense, AI has been described as “the study of the computations that make it possible to perceive, reason, and act”³⁰ or “the automation of intelligent behavior”,³¹

24 Ibid.

25 For an overview of various definitions, and an attempt to formalize a concept of (machine) intelligence in the broadest reasonable sense, see Shane Legg and Marcus Hutter, “Universal Intelligence: A Definition of Machine Intelligence,” *Minds and Machines* 17, no. 4 (2007): 391–444.

26 Cf. Yudkowsky, “Artificial Intelligence as a Positive and Negative Factor in Global Risk.”, p. 311

27 Shane Legg and Marcus Hutter, “A Collection of Definitions of Intelligence,” *arXiv:0706.3639 [Cs]*, June 25, 2007, <http://arxiv.org/abs/0706.3639>.

28 Ibid.

29 John McCarthy and Stanford University Formal Reasoning Group, “What Is Artificial Intelligence | Basic Questions,” *Formal Reasoning Group*, 2007, <http://www-formal.stanford.edu/jmc/whatisai/node1.html>.

30 Patrick Henry Winston, *Artificial Intelligence*, 3rd ed., 1992.

31 George F. Luger and William A. Stubblefield, *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*, 6th ed., 2008. ; cf. also A. Barr and Feigenbaum, eds., *The Handbook of Artificial Intelligence*, vol. 2 [Stanford,

which is driven by a general “study of intelligent agents” both biological and artificial.³² There are furthermore dozens of definitions and typologies of what constitutes artificial intelligence.³³ However, in concrete terms, and in most applications, AI is defined as non-human intelligence that is measured by its ability to replicate human mental skills, such as pattern recognition, understanding natural language (NLP), adaptive learning from experience, strategizing, or reasoning about others.³⁴

Likewise, militaries have also considered AI in a functional context – by the degree to which a machine, in the words of one (older) study of the US Army Sciences Board, “... can incorporate abstraction and interpretation into information processing and make decisions at a level of sophistication that would be considered intelligent in humans”,³⁵ a definition that is also upheld in the more recent Summer Study on Autonomy by the US Defense Science Board, which describes AI as “the capability of computer systems to perform tasks that normally require human intelligence (e.g., perception, conversation, decision-making).”³⁶ These characterizations, and the explicit link to human performance, derive from the very inception of the field of AI. The original proposal for the seminal 1956 Dartmouth College Summer Project argued that:

*“Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make **machines that use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.**”³⁷*

However, while human intelligence has been the common choice as a yardstick for benchmarking or assessing progress in AI development, there are also approaches that do not seek to recreate human intelligence or performance, but instead focus more on systems that approach an ideal-typical ‘rational’ performance. As a result, most concrete definitions of AI fall into one of four categories (cf. **Figure 3**); these categories represent approaches which take distinct positions on two conceptual dimensions:

1. Whether it emphasises the attainment of specific (‘intelligent’ or ‘sentient’) **thought processes (T)** and reasoning, or whether it emphasises (‘goal-oriented’; ‘effective’) **behavior (B)**;
2. Whether it measures success on (1) against **human performance (H)**, or against an ideal concept of intelligence – usually defined as **‘rationality’ (R)**;³⁸

California & Los Altos, California: HeurisTech Press and William Kaufmann, 1982).

32 Cf. Poole, et al., 1998, p.1; Russell, & Norvig, 2009

33 For a comprehensive overview, see also Daniel Faggella, “What Is Artificial Intelligence? An Informed Definition -,” *TechEmergence.com*, October 10, 2016, <http://techemergence.com/what-is-artificial-intelligence/>

34 Cf. Russell, & Norvig, 2009, one of the standard textbooks in the field.

35 Barry J. Brownstein and et al., “Technological Assessment of Future Battlefield Robotic Applications,” Proceedings of the Army Conference on Application of AI to Battlefield Information Management (White Oak: US Navy Surface Weapons Center, 1983), p. 169

36 Defense Science Board, “Report of the Defense Science Board Summer Study on Autonomy” (Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, June 2016), <https://www.hsdl.org/?view&did=794641>, p. 5

37 J. McCarthy et al., “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” August 31, 1955, <http://robotics.cs.tamu.edu/dshell/cs625/2014-09-02.pdf>. Emphasis added.

38 Russell and Norvig, *Artificial Intelligence: A Modern Approach*

	Human Benchmark (H)	Rationality benchmark (R)
Intelligence as Thought Processes (T)	<p>(T-H) Systems that think like humans (e.g. cognitive science)</p> <p><i>“The exciting new effort to make computers think ... machines with minds, in the full and literal sense”</i> Haugeland, 1985</p> <p><i>“The automation of activities that we associate with human thinking, activities such as decision-making, problem solving, learning ...”</i> Bellman, 1978</p>	<p>(T-R) Systems that think rationally (logic/laws of thought)</p> <p><i>“The study of mental faculties through the use of computational models”</i> Charniak and McDermott, 1985</p> <p><i>“The study of the computations that make it possible to perceive, reason, and act”</i> Winston, 1992</p>
Intelligence as goal-oriented behavior (B)	<p>(B-H) Systems that act like humans (Cf. Turing test; Winograd Schema Challenge³⁹)</p> <p><i>“The art of creating machines that perform functions that require intelligence when performed by people”</i> Kurzweil, 1990</p> <p><i>“The study of how to make computers do things at which, at the moment, people are better”</i> Rich and Knight, 1991</p>	<p>(B-R) Systems that act rationally (rational agents)</p> <p><i>“A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes”</i> Schalkoff, 1990</p> <p><i>“The branch of computer science that is concerned with the automation of intelligent behavior”</i> Luger & Stubblefield, 1993</p>

FIGURE 3: A TYPOLOGY OF DIFFERENT DEFINITIONS OF AI, AND THEIR UNDERLYING APPROACHES, WITH SOME EXAMPLES³⁹

Ultimately, while approaches that emphasize internal (e.g. ‘T’) models are of intellectual and philosophical interest, it is important that the definition of AI which we use does not put arbitrary limits on the inner workings of AI, or the approaches used to creating it.⁴⁰ Moreover, from a practical (that is, political or military) perspective, the focus tends to be on the way artificially intelligent systems manage to *act* in the world, that is behavior-focused (B) approaches. Secondly, both in current ‘deep learning’ applications, as well as in most prospective military applications, AI systems will provide the most added value (the competitive edge), not only by equating human intelligence, but precisely by surpassing it – even if only within a narrow domain, such as information analysis or reaction time.

³⁹ Adapted from Russell and Norvig, *Artificial Intelligence: A Modern Approach*, p. 5. The sources mentioned are R.E. Bellman, *An Introduction to Artificial Intelligence: Can Computers Think?* (San Francisco: Boyd & Fraser Publishing Company, 1978).; E. Charniak and D. McDermott, *Introduction to Artificial Intelligence* (Massachusetts: Addison-Wesley Reading, 1985).; J. Haugeland, ed., *Artificial Intelligence: The Very Idea* (Cambridge, Mass: MIT Press, 1985).; R. Kurzweil, *The Age of Intelligent Machines* (Cambridge, Mass: MIT Press, 1990).; George F. Luger and William A. Stubblefield, *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*, 2nd ed. (California: Benjamin/Cummings, 1993).; E. Rich and K. Knight, *Artificial Intelligence*, 2nd ed. (New York: McGraw-Hill, 1991).; R.I. Schalkoff, *Artificial Intelligence: An Engineering Approach* (New York: McGraw-Hill, 1990). and Winston, *Artificial Intelligence*.

⁴⁰ Cf. Shane Legg and Marcus Hutter, “A Formal Definition of Intelligence for Artificial Systems,” *Proc. 50th Anniversary Summit of Artificial Intelligence*, 2006, 197–198.; cf. Faggella, “What Is Artificial Intelligence?”

Accordingly, one strategically (and therefore analytically) relevant definition for our purposes is that of AI as a rational optimization agent that can (enable humans to) act competently in the world (that is, subtype B-R). Therefore, while it is important to use human performance as one measuring stick to understand the threshold at which AI systems (or so-called ‘centaur’ human-AI teams) can start to outperform and outcompete human adversaries, for most strategic purposes we can focus on the thought processes (T), and specifically the behavioral performance (B) of AI systems, at an equal- or better-than human level of accuracy, speed, or decision quality. This corresponds to the definition provided by Nils J. Nilsson; “Artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.”⁴¹

2.1.3. Artificial Narrow-, General-, and Superintelligence

There is one additional distinction that is of particular use when thinking about the sequence and speed with which progress on AI will unfold. The literature generally agrees that there are three tiers of AI⁴² – which can also be seen as three consecutive generations of AI (see also chapter 3.6 page 55):

- **Artificial Narrow Intelligence** (ANI or “narrow AI”)⁴³: machine intelligence that *equals or exceeds* human intelligence *for specific tasks*. Existing examples of such systems would be IBM’s Deep Blue (Chess) & Watson (‘Jeopardy!’), Google’s AlphaGo (go), High-Frequency Trading Algorithms, or indeed any specialized automatic systems performing beyond human reach (e.g. Google Translate; spam filters; the guidance systems of point-defense anti-missile cannons etc.).
- **Artificial General Intelligence** (AGI or “strong AI”): machine intelligence meeting the *full range* of human performance across any task; and:
- **Artificial Superintelligence** (ASI): machine intelligence that *exceeds human intelligence* across any task.⁴⁴

41 Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (Cambridge ; New York: Cambridge University Press, 2010).

42 Another frequently used classification differentiates between weak and strong AI. There seem to be quite a few different ways to characterize the difference (for a divergent one, see Stephen Lucci and Danny Kopec, *Artificial Intelligence in the 21st Century: A Living Introduction*, 2016. but the most frequently used one differentiates between “weak AI”—the variety devoted to providing aids to human thought—and “strong AI”—the variety that attempts to mechanize human-level intelligence (a machine with consciousness, sentience and mind). See Nilsson, *The Quest for Artificial Intelligence*

43 Ray Kurzweil, “Long Live AI,” *Forbes.com*, August 15, 2015, http://www.forbes.com/home/free_forbes/2005/0815/030.html

44 Note that even such hypothetical superintelligences might come in different forms, such as ‘speed superintelligence’ - ‘a system that can do all that a human intellect can do, but much faster’; ‘collective superintelligence’ - ‘A system composed of a large number of smaller intellects such that the system’s overall performance across many very general domains vastly outstrips that of any current cognitive system’; or ‘quality superintelligence’ - “A system that is at least as fast as a human mind and vastly qualitatively smarter.” Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (OUP Oxford, 2014),. pg. 53, 54, 56

2.2. A Brief History of AI

Contrary to widespread perception AI is not an entirely new discipline. Many of its founding concepts draw on over 2000 years of insights accumulated in philosophy, logic, mathematics, theories of reasoning, cognitive psychology and linguistics.⁴⁵ As a practical, applied field, however, AI truly came into its own in the direct wake of the Second World War. The inception of AI as a discipline occurred in wartime research in areas such as cryptography and the calculation of ballistic firing tables for artillery,⁴⁶ and it was spurred on by the foundational work of Alan Turing,⁴⁷ as well as the work on simple neural networks by Warren McCulloch and Walter Pitts.⁴⁸ AI first became a distinct field of study during the small Dartmouth Summer Project held in 1956 – an event that was to herald the first ‘spring’ for AI research.

2.2.1. Early Enthusiasm: The First AI Spring (‘56-75’)

Contrary to the ambitious aspirations of the Dartmouth project – its ten attendees had anticipated making “a significant advance” within the space of a single summer⁴⁹ – the workshop did not lead to any direct breakthroughs. However, it did give rise to the name ‘artificial intelligence’, and by introducing all major thinkers to each other, it established the shape of that field for decades to come.⁵⁰ The intellectual cross-pollination at Dartmouth, and a subsequent string of high-profile achievements by AI researchers in developing tools capable of human-like performance in narrow fields such as geometrical proofs, algebra, and simple games (e.g. checkers⁵¹), led to an early golden age of AI research (1956-1974). This was a period in which prototype systems – operating safely within well-defined limits of an artificially simple ‘microworld’ – again and again refuted established beliefs that *‘no machine could ever do X!’*⁵²

In addition to systems based on formal symbolic reasoning, there was a line of ‘connectionist’ research into what became known as *perceptrons*, an early form of neural networks.⁵³ During this period, considerable amounts of scientific – and, with the increasing pressures of the Cold War, also military – funds were directed into the field.⁵⁴ From the early 1960s to the 1970s, and following a philosophy that it should ‘fund people,

45 Russell and Norvig, *Artificial Intelligence: A Modern Approach.*, p. 8-15

46 In the shape of ENIAC, the first electronic general-purpose computer, which was developed from 1943-1945 (though not delivered until the end of the war) by John Mauchly and J. Presper Eckert. Cf. Harry L. Reed Jr., “Firing Table Computations on the Eniac,” in *Proceedings of the 1952 ACM National Meeting (Pittsburgh)*, ACM ‘52 (New York, NY, USA: ACM, 1952), 103–106, doi:10.1145/609784.609796

47 e.g. A.M. Turing, “Computing Machinery and Intelligence,” *Mind* 59, no. 236 (October 1950): 433–60

48 Warren S. McCulloch and Walter Pitts, “A Logical Calculus of the Ideas Immanent in Nervous Activity,” *Bulletin of Mathematical Biophysics* 5 (1943): 115–33

49 McCarthy et al., “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence.”

50 Russell and Norvig, *Artificial Intelligence: A Modern Approach.*, pg. 17

51 Cf. Arthur L. Samuel, “Some Studies in Machine Learning Using the Game of Checkers,” *IBM Journal* 3, no. 3 (July 1959), <https://www.cs.virginia.edu/~evans/greatworks/samuel1959.pdf>.

52 Cf. Bostrom, *Superintelligence.*, pg. 5

53 Russell and Norvig, *Artificial Intelligence: A Modern Approach.*, pg. 19-20

54 Daniel Crevier, *AI: The Tumultuous Search for Artificial Intelligence* (New York: BasicBooks, 1993).; Hans Moravec, *Mind Children: The Future of Robot and Human Intelligence*, 4. print (Cambridge: Harvard Univ. Press, 1995),pg. 9

not projects’,⁵⁵ the Advanced Research Projects Agency (ARPA) was to provide major funding, with few strings attached, to open-ended AI research projects of renowned researchers at MIT, Carnegie Mellon University, and Stanford, amongst others.

2.2.2. ‘Grandiose Objectives’: The First AI Winter (‘74-’80)

This first wave of enthusiasm proved unsustainable, however, and by the early 1970s the field of AI was slipping into the first AI Winter. It became clear that rapid early progress had slowed down, and that AI systems remained far more limited in their capabilities than had initially been expected. There were several reasons for this disillusion, chief amongst them the discovery of the combinatorial explosion of possibilities’, which demonstrated that exhaustive search, a key trial-and-error approach underpinning many AI algorithms at the time, would require exponential amounts of computer processing power once one moved away from ‘toy problems’ to more complex real-world problems – an amount of processing power which was unimaginable at the time. This challenge was aggravated by the technical limitations at the time – chiefly hardware limits on memory and processor speed.⁵⁶

There were also significant internal disagreements; in a seminal 1969 book,⁵⁷ Marvin Minsky and Seymour Papert demonstrated that there were limits to the performance of the early perceptrons of the connectionist school. As a result of this work, research funding on these early neural nets was mostly abandoned for a decade, and work focused more on programming formal rules into symbolic systems, which would feed into the later expert system’ of the 1980s.

All these diverse problems led to significant funding cutbacks on both sides of the Atlantic; in 1966, a scathing report by the Automatic Language Processing Advisory Committee of the US government on the slow progress in natural language processing and automatic translation led to the termination of all research funding by the US National Research Council.⁵⁸ In 1973, the publication of the Lighthill Report⁵⁹, which was highly critical of the failure of AI research to live up to its ‘grandiose objectives’, led to a complete cessation of British research for over a decade, sending a bow wave of funding cuts in AI research across Europe. At the same time, ARPA was deeply disappointed in the returns on its initial investment in AI, specifically in speech recognition, but also more broadly. After the passage of the 1969 Mansfield Amendment, which required US military research funding to focus more on “mission-oriented direct research, rather than basic undirected research”⁶⁰, it did not take long for nearly all military funding to dry up.

55 Daniel Crevier, *AI: The Tumultuous Search for Artificial Intelligence* (New York: BasicBooks, 1993). p. 65

56 Bostrom, *Superintelligence.*, pg. 6-7

57 Marvin Minsky and Seymour Papert, *Perceptrons* (Oxford, England: M.I.T. Press, 1969).

58 John Hutchins, “The History of Machine Translation in a Nutshell,” 2005. <http://www.hutchinsweb.me.uk/Nutshell-2005.pdf>.

59 James Lighthill, “Artificial Intelligence: A General Survey,” *Artificial Intelligence: A Paper Symposium* (Science Research Council, 1973), <http://www.math.snu.ac.kr/~hicho/infomath/Articles/Lighthill%20Report.pdf>.

60 The Mansfield Amendment was passed as part of the Defense Authorization Act of 1970. See US Congress, “Defense Authorization Act,” Pub. L. No. 91–121 (1969).. See also Committee on Innovations in Computing and Communications, National Research Council, “Chapter 9: Developments in Artificial Intelligence | Funding a

2.2.3. Expert Systems: The Second AI Spring ('80-'87)

A second wave came in the 1980s, with the advent of so-called expert systems – rule-based programmes that answered questions or solved problems within a small domain of specific knowledge.

Such programmes emulated decision-making processes of a human expert, which had been hard-coded into a formal logical language. Expert systems could be used as support tools for decision-makers or executives. At the same time, after almost a decade of relative standstill, research on neural networks, the aforementioned connectionist school, saw a renaissance spurred on by progress in physics and computer sciences, as well as the development of the backpropagation algorithm,⁶¹ which made it possible to train multi-layered neural networks capable of learning a wider range of useful functions.

Meanwhile, the promise of ANI-like expert systems – applications of AI which were recognized as economically useful by the broader industry⁶² – led to interest and funding for research, starting with Japan's launch (in 1981) of the 10-year *Fifth-Generation Computer Systems Project*, a public-private partnership that sought to develop a massively parallel computing architecture capable of rapidly drawing on vast stores of rules, producing AI capable of natural conversation, translation, and image interpretation.⁶³ Spurred on by fears of falling behind, many Western governments followed suit and provided funding to AI research – such as the Microelectronics and Computer Technology Corporation (MCC) in the US, and the Alvey Project in the UK. Simultaneously, DARPA founded its Strategic Computing Initiative in the process tripling its funding for research.⁶⁴ As a result, various software and hardware companies converged around expert systems and sales went from a few millions in 1980 to \$2 billion by 1988.⁶⁵

2.2.4. 'Not the Next Wave'; The Second AI Winter ('87-'93)

Once again, however, the boom turned into a bust, as research programs in Japan, the US and Europe failed to meet their objectives. Many specialized AI hardware companies collapsed after 1987, as the desktop computers from Apple and IBM rapidly outcompeted more expensive AI-specific ones. With only some exceptions, expert systems proved of limited practical utility. Small expert systems generally were ineffective or added little

Revolution: Government Support for Computing Research," in *Funding a Revolution: Government Support for Computing Research*, 2008, <https://web.archive.org/web/20080112001018/http://www.nap.edu/readingroom/books/far/ch9.html>

61 The first algorithm underlying backpropagation had already been discovered in 1969, by Arthur Bryson and Yu-Chi Ho, Arthur Earl Bryson and Yu-Chi Ho, *Applied Optimal Control: Optimization, Estimation, and Control*. (Blaisdell Publishing Company or Xerox College Publishing, 1969).. Its application to neural nets was developed by Paul Werbos, Paul Werbos, "Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences" (Ph.D., Harvard University, 1974).. However the algorithm was only popularized in the community in the mid-1980s by David Rummelhart, in David E. Rumelhart, James L. McClelland, and CORPORATE PDP Research Group, eds., *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations* (Cambridge, MA, USA: MIT Press, 1986)

62 Russell and Norvig, *Artificial Intelligence: A Modern Approach*., p23-24

63 Ibid., p24; Bostrom, *Superintelligence*., p7

64 Pamela McCorduck, *Machines Who Think*, 2nd ed. (Natick, MA: A.K. Peters, Ltd., 2004).

65 Russell and Norvig, *Artificial Intelligence: A Modern Approach*.

value, large ones could be prohibitively expensive to create, test and keep updated and they were also fairly 'brittle', as they tended to crash in the face of unusual inputs.⁶⁶ At the end of the 1980s, the leadership at DARPA's Information Processing Technology Office took the view that AI was not "the next wave" and that researchers had once again over-promised and under-delivered.

This led to severe funding cuts across the Strategic Computing Initiative, leaving only programs of direct military relevance – including an autonomous battle tank program, never delivered, and the DART logistics & battle management system, which would later prove invaluable during the First Gulf War.⁶⁷

2.2.5. A Third, Sustained AI Spring ('93-'11)

After the relative disappointment of the previous decades, many AI researchers abandoned long-term dreams of developing 'human-level' AI, and instead moved into fragmented sub-fields focused on rigorously solving specific problems or applications.⁶⁸ Moreover, the traditional logic-based paradigm, which had reached its summit with the imposing but brittle expert systems of the 1980s, found itself progressively challenged by techniques both old and new. Perhaps the most important of these came in the form of a reinvigorated interest in neural networks and genetic algorithms – approaches which distinguished themselves from the earlier expert systems by exhibiting 'graceful degradation' – referring to small errors in the assumptions resulting in only small reductions in performance, rather than a full crash.

As a result of this pragmatic turn, further fuelled by Moore's Law and the attendant advances in hardware capabilities,⁶⁹ the field of AI quietly began to bloom.⁷⁰ It began not only to achieve some of its decades-old goals, but was also able to increasingly, if invisibly, permeate both the technology industry, and, frequently, aspects of daily life.⁷¹ In the form of direct applications or downstream spinoff technologies, research was finding utility in a wide range of fields, from games (notably the famous 1997 chess victory, of IBM's Deep Blue over world champion Garry Kasparov); to logistics, spacecraft and satellite monitoring; robotics; traffic management; medical diagnostics;⁷² speech recognition,⁷³ autonomous vehicles and Google's search engines, to name but a few.⁷⁴

66 Douglas Lenat and R.V. Guha, *Building Large Knowledge-Based Systems* (Addison-Wesley, 1989).. See also the account in Stuart Russell, "Defining Intelligence," *EDGE*, February 7, 2017, <https://www.edge.org/conversation/stuart-russell-defining-intelligence>

67 McCorduck, *Machines Who Think.*, pg. 430-43

68 Ibid

69 Cf. Ray Kurzweil, *The Singularity Is Near* (Viking Press, 2005)., p274.


70 Cf. Jiqiang Niu et al., "Global Research on Artificial Intelligence from 1990–2014: Spatially-Explicit Bibliometric Analysis," *ISPRS International Journal of Geo-Information* 5, no. 5 (May 16, 2016): 66, doi:10.3390/ijgi5050066

71 Russell and Norvig, *Artificial Intelligence: A Modern Approach*.

72 Ibid., pp. 26-27;

73 Cf. The Economist, "Are You Talking to Me? -," *The Economist*, June 7, 2007, <http://www.economist.com/node/9249338>.

74 Ajit Nazre and Rahul Garg, "A Deep Dive in the Venture Landscape of Artificial Intelligence and Machine Learning," September 2015, <http://slideplayer.com/slide/7002258/>.



Military applications incorporating ANI also grew in prominence such as for instance unmanned aerial vehicles (UAVs). Military UAVs had previously seen limited use, by the US in the Vietnam War, and by the Israeli Defense Force during the Yom Kippur War ('73) and (with great success) against the Syrian air force during the First Lebanon War ('82).⁷⁵ However, from the '90s, and with the first deployment of the Predator UAV during the Balkan Wars, these systems started to become a widely used component of military operations – and UAVs have made appearances in almost every conflict since.

Finally, the performance of the previously mentioned DART tool for automated logistics planning and scheduling, development of which had survived even the drawdown in DARPA's Strategic Computing Initiative, during Operation Desert Storm in 1991, led DARPA to claim that this single application had more than paid back their thirty-year investment in AI.⁷⁶

In addition to private funding, public funding picked up as well: since the mid-2000s, both the European Union (under its Seventh Framework Programme) and the US (through DARPA's Grand Challenge Program and the Cognitive Technology Threat Warning System) have begun to provide renewed broad funding for AI research. As such, even though it took until the mid-2000s for AI researchers to shrug off any negative legacies associated with the term AI – many instead choosing to describe their work as 'knowledge-based systems'; 'informatics', or 'computational intelligence', to name a few⁷⁷ – Research gradually but surely has entered its third, and so far most sustained 'spring' a steady blooming.

2.2.6. Big Data, Deep Learning, and an Artificial Intelligence Revolution ('11-present)

In recent years, the '3rd reincarnation'⁷⁸ of AI appears to have hit a tipping point. This success is driven by specific factors which have produced a dramatic improvement in the predictive accuracy of algorithms. Some of these factors are conceptual, including advances in neuroscience as well as computer science – most notably the work of Geoffrey Hinton and Salakhutdinov⁷⁹, which pioneered powerful new techniques for enabling neural network pattern-recognition. Other factors of note are increases in affordable computing power and faster networks; cloud infrastructures; the growth in the Internet of Things and Big Data – and specifically the open-source availability of very large datasets (sometimes

75 Mary Dobbing and Chris Cole, "Israel and the Drone Wars: Examining Israel's Production, Use and Proliferation of UAVs" (Drone Wars UK, January 2014), <https://dronewarsuk.files.wordpress.com/2014/01/israel-and-the-drone-wars.pdf>.

76 Sara Reese Hedberg, "DART: Revolutionizing Logistics Planning," *IEEE Intelligent Systems* 17, no. 3 (May 2002): 81–83, doi:10.1109/MIS.2002.1005635.; Stephen E. Cross and Edward Walker, "DART: Applying Knowledge Based Planning and Scheduling to CRISIS Action Planning.," in *Intelligent Scheduling*, ed. Monte Zweben and Mark Fox (San Francisco, CA: Morgan Kaufmann, 1994), 711–29.

77 John Markoff, "Behind Artificial Intelligence, a Squadron of Bright Real People," *The New York Times*, October 14, 2005, <http://www.nytimes.com/2005/10/14/technology/behind-artificial-intelligence-a-squadron-of-bright-real-people.html>.

78 Nazre and Garg, "A Deep Dive in the Venture Landscape of Artificial Intelligence and Machine Learning.," slide 2

79 G. E. Hinton and R. R. Salakhutdinov, "Reducing the Dimensionality of Data with Neural Networks," *Science* 313, no. 5786 (July 28, 2006): 504–7, doi:10.1126/science.1127647.

generated from social-media networks) for use in the training and testing of large-scale machine learning networks.⁸⁰

Furthermore, more than ever before, AI research has benefited tremendously from a huge boom in the level, diversity and sources of funding and talent, including from major private sector players such as Apple, Amazon, Baidu, Google, Facebook, IBM, Microsoft.⁸¹ For many of these, AI systems increasingly are not just one more asset by which to marginally increase existing profits: instead, they are at the heart of the enterprise's business model. AI research has also benefited from investments by traditional top-spenders on research, amongst others from the automotive industry – with \$1 billion investment programs by Toyota⁸² and (in February 2017) Ford Motors,⁸³ alongside major programs by Mercedes-Benz,⁸⁴ BMW,⁸⁵ – as well as the pharmaceutical sector.⁸⁶

Finally, all this has also been facilitated by a new mode of cumulative knowledge-building, which marks a shift away from closed epistemic communities with limited interconnections that are tightly policed by gatekeepers, towards an open-source model of research, debate, scholarship and content creation (as in the case of Google Translate, an application which, calibrated with reference to existing document translations, and continually improved by massed user feedback, has rapidly bypassed the accuracy of any of the traditional translation programs, which relied on painstaking manual coding by linguistic experts).

Whatever other effects (good and bad) this shift towards collaborative ecosystem-like crowdsourcing has had, it has certainly lowered barriers to access, created greater (reputational and financial) rewards for success and innovation, and led to the proliferation of ever more comprehensive and widely available data, which in turn has allowed for the rapid implementation, deployment, testing and iteration over new AI applications.⁸⁷

80 Nazre and Garg, "A Deep Dive in the Venture Landscape of Artificial Intelligence and Machine Learning."

81 Ibid., slide 7

82 Rachel Metz, "Toyota Investing \$50M for Autonomous-Car Research at Stanford and MIT," *MIT Technology Review*, 2015, <https://www.technologyreview.com/s/541046/toyota-investing-50m-with-stanford-mit-for-autonomous-car-research/>; Hans Greimel, "What's behind Toyota's Big Bet on Artificial Intelligence?," *Automotive News*, 2015, <http://www.autonews.com/article/20151109/OEM06/311099937/whats-behind-toyotas-big-bet-on-artificial-intelligence%3F>.

83 Mike Isaac and Neal E. Boudette, "Ford to Invest \$1 Billion in Artificial Intelligence Start-Up," *The New York Times*, February 10, 2017, <https://www.nytimes.com/2017/02/10/technology/ford-invests-billion-artificial-intelligence.html>

84 Mark Bergen, "Mercedes-Benz Wants to Beat Google, Uber to Our Driverless Future," *Recode*, November 26, 2015, <http://www.recode.net/2015/11/26/11620962/mercedes-benz-wants-to-beat-google-uber-to-our-driverless-future>

85 Horatiu Boeriu, "BMW Sees New Opportunities in Artificial Intelligence," *BMW BLOG*, June 9, 2016, <http://www.bmwblog.com/2016/06/09/bmw-sees-new-opportunities-artificial-intelligence/>.

86 Cf. CB Insights, "From Virtual Nurses To Drug Discovery: 106 Artificial Intelligence Startups In Healthcare," *CB Insights - Blog*, February 3, 2017, <https://www.cbinsights.com/blog/artificial-intelligence-startups-healthcare/>, for an description of startups applying AI to medical applications, see João Medeiros, "The Startup Fighting Cancer with AI," *WIRED UK*, accessed February 25, 2017, <http://www.wired.co.uk/article/ai-cancer-drugs-berg-pharma-startup..> For one overall overview of AI investments by sector, see Bank of America Merrill Lynch, "Robot Revolution - Global Robot & AI Primer," Primer, Thematic Investing [Bank of America Merrill Lynch, December 2015], http://www.bofam.com/content/dam/boamlimages/documents/PDFs/robotics_and_ai_condensed_primer.pdf.

87 Stanford University, "Artificial Intelligence and Life in 2030 One Hundred Year Study on Artificial Intelligence | Report of the 2015 Study Pane.", p. 14-17

This potent combination of powerful hardware, extensive funding, (relatively) open development, and big (labelled) test datasets, has resulted in remarkable achievements in numerous fields. This includes the long-intractable field of natural language processing: starting with IBM Watson's famous 2011 Jeopardy! victory, this track record has continued with the public debut of voice-responsive Virtual Personal Assistants (like Apple's SIRI, or Microsoft's Cortana), culminating in the 2015 DeepMind project which trained a deep neural net on over 300,000 CNN and Daily Mail articles, after which it was able to use information from these articles to accurately answer 60% of the queries put to it.⁸⁸ Other innovations marked the first challenges to the hallowed Turing test, including the 2013 announcement, by startup Vicarious announcing in 2013 that its AI had achieved a 90% success rate at solving CAPTCHA tests;⁸⁹ and the development of 'Eugene Goostman', a chatbot which in 2014 managed to fool 33% of human judges at the Royal Society that it was a 13-year old boy.⁹⁰

Other advances have come in the field of face recognition, with various algorithms developed by Google (GoogLeNet), Facebook (DeepFace – 2013) and Yahoo (DeepDense – 2015)⁹¹ having achieved better-than-human identification rates. Finally, Google DeepMind has achieved remarkable successes in developing self-teaching AI agents capable of superhuman performance at common video games (ATARI Deep-Q – 2013)⁹² and the game of go (AlphaGo – 2016).⁹³

There has also been growing interest in the (cyber)security applications of these technologies; in the Summer of 2016, the 'Mayhem' AI (Carnegie Mellon University) won the DARPA Cyber Grand Challenge in automatic system vulnerability detection & patching.⁹⁴

88 Emerging Technology from the MIT Tech Review, "Google DeepMind Teaches Artificial Intelligence Machines to Read," *MIT Technology Review*, 2015, <https://www.technologyreview.com/s/538616/google-deepmind-teaches-artificial-intelligence-machines-to-read/>. See also the original paper at Karl Moritz Hermann et al., "Teaching Machines to Read and Comprehend," *arXiv:1506.03340 [Cs]*, June 10, 2015, <http://arxiv.org/abs/1506.03340>.

89 Rachel Metz, "AI Startup Vicarious Claims Victory Over Captchas," *MIT Technology Review*, 2013, <https://www.technologyreview.com/s/520581/ai-startup-says-it-has-defeated-captchas/>.

90 The Guardian, "Computer Simulating 13-Year-Old Boy Becomes First to Pass Turing Test," *The Guardian*, June 9, 2014, sec. Technology, <https://www.theguardian.com/technology/2014/jun/08/super-computer-simulates-13-year-old-boy-passes-turing-test>; others however questioned the significance of the contest, cf. Gary Marcus, "What Comes After the Turing Test?," *The New Yorker*, June 9, 2014, <http://www.newyorker.com/tech/elements/what-comes-after-the-turing-test>

91 Cf. Emerging Technology from the MIT Tech Review, "The Face Detection Algorithm Set to Revolutionize Image Search," *MIT Technology Review*, 2015, <https://www.technologyreview.com/s/535201/the-face-detection-algorithm-set-to-revolutionize-image-search/>; see also the original paper at Sachin Sudhakar Farfade, Mohammad Saberian, and Li-Jia Li, "Multi-View Face Detection Using Deep Convolutional Neural Networks," *arXiv:1502.02766 [Cs]*, February 9, 2015, <http://arxiv.org/abs/1502.02766>.

92 Volodymyr Mnih et al., "Playing Atari with Deep Reinforcement Learning," 2013.; Cade Metz, "Teaching Computers to Play Atari Is A Big Step Toward Bringing Robots Into the Real World," *WIRED*, 2015, <https://www.wired.com/2015/12/teaching-ai-to-play-atari-will-help-robots-make-sense-of-our-world/>

93 In February 2017, another AI company, Libratas, also announced a breakthrough in AI systems playing poker against the best human players. cf. Cade Metz, "Inside the Poker AI That Out-Bluffed the Best Humans," *WIRED*, 2017, <https://www.wired.com/2017/02/libratas/>.

94 Nonetheless, Mayhem was still soundly defeated by all-human hacking teams at the DEFCON summit soon after. See Devin Coldewey, "Carnegie Mellon's Mayhem AI Takes Home \$2 Million from DARPA's Cyber Grand Challenge," *TechCrunch*, 2016, <http://social.techcrunch.com/2016/08/05/carnegie-mellons-mayhem-ai-takes-home-2-million-from-darpas-cyber-grand-challenge/>

Finally, there has also been a growing awareness of the societal and operational concerns revolving around the accountability of ANI systems which are based on opaque machine learning capabilities. For instance, in the fall of 2016, DARPA issued a funding call for projects focused on 'Explainable Artificial Intelligence' (XAI) – “new or modified machine learning techniques that produce explainable models that, when combined with effective explanation techniques, enable end users to understand, appropriately trust, and effectively manage the emerging generation of Artificial Intelligence (AI) systems.”⁹⁵



The Tabulating Era
(1900s–1940s)

The Programming Era
(1950s–present)

The Cognitive Era
(2011–)


FIGURE 4: IBM’S VIEW OF THE HISTORY OF COMPUTING⁹⁶

Recent developments strongly suggest that AI research is now moving into the third of three main ‘eras’ in the evolution of AI, identified by IBM’s John Kelly.⁹⁷ In the first era (the ‘Tabulating Era’ – 1900-1940s), humans created single-purpose mechanical systems that essentially were able to perform a single function: to count. They used punched cards to input and store data, and to eventually instruct the machine what to do – in a very primitive way. In the second era (the ‘Programming Era’ – 1950s-present), the advent of electronic systems allowed for the evolution of digital “computers” that were able to perform ever more sophisticated if/then logical operations and loops, with instructions coded in software that was programmed by humans. With the rise of powerful deep learning systems, and their application throughout research and industry, we are at last entering the Cognitive Era, in which these software systems no longer have to be fully programmed, but can instead start ‘learning’ autonomously through training, use,

95 DARPA, “Grant Opportunity - DARPA-BAA-16-53 - Explainable Artificial Intelligence (Xai)” (Department of Defense - DARPA - Information Innovation Office, n.d.), <http://www.grants.gov/web/grants/view-opportunity.html?oppld=287284>. See also David Gunning, “Explainable Artificial Intelligence (XAI),” DARPA, accessed September 23, 2016, <http://www.darpa.mil/program/explainable-artificial-intelligence>.

96 J. E. Kelly, “Computing, Cognition and the Future of Knowing,” *Whitepaper, IBM Research*, 2015, http://www.academia.edu/download/44915350/Computing_Cognition_WhitePaper.pdf

97 Ibid



and user feedback, in ways which may resemble human learning, but which often can no longer be fully comprehended by it.⁹⁸ To understand the opportunities – and limits – offered by such learning machines, we will now briefly review the AI discipline, and its distinct approaches at developing machines that can accurately represent the world, learn, and take appropriate, and increasingly intelligent decisions.

2.3. AI: A Cookbook of Components, Approaches and Design Architectures

For those unfamiliar with the field, or too familiar with popular culture, it may sometimes appear as if the development of Artificial Intelligence, insofar as it seeks to reproduce human intelligence (singular), is about a single technology – a unified search to develop that kernel of sentience in machines, which humans already possess. Yet nothing could be further from the truth: in practice, far from involving just one single technology or unified discipline, AI is a collection of distinct fields working on different technologies that, working together in the appropriate environments, allow for intelligent behavior to emerge.

2.3.1. Design Specifications: The Components of Intelligence

In this context, it helps to recall that the human variety of intelligence is similarly complex and multifaceted. When we think of human intelligence, we tend to think of a singular abstract cognitive ability that we see as a trait of an individual, making her ‘smarter’ (e.g. more knowledgeable; more capable) than others. When we do so, however, we tend to forget how many different – (neuro)biological – components have to coalesce in order for any form of intelligence to materialize. Let us return to the example of intelligence we started with: reading. As you read these sentences, the center of the retinas in your eyes has to use motor movement to scan the page carefully. This scan is then transferred as chemical signals to the visual cortex in the brain which recognizes patterns in a row of black marks against a white background. Your visual system then progressively extracts graphemes, syllables, prefixes, suffixes, and word roots as it parses our natural language. Through various routes – which again involve various parts of our brain – it converts those few black marks into an entire semantic universe of meanings by cross-referencing them against various other pieces of information stored in your brain⁹⁹. And all of this is just the ‘legere’ part of ‘intelligence’ – the overall process whereby reading (or hearing, or talking, or experiencing, etc.) translates into intelligence that is even more complex¹⁰⁰.

98 For a thought-provoking analysis of the ‘inscrutability’ of deep-learning based AI networks, see Aaron M. Bornstein, “Is Artificial Intelligence Permanently Inscrutable?,” *Nautilus*, 2016, <http://nautilus.us/issue/40/learning/is-artificial-intelligence-permanently-inscrutable>.

99 For a fascinating overview of this process, see Stanislas Dehaene, *Reading in the Brain: The Science and Evolution of a Human Invention* (New York: Viking, 2009). For a more recent overview of the different parts of the brain that are known to be involved in this process, see Stephani Sutherland, “What Happens in the Brain When We Read?,” *Scientific American Mind* 26, no. 4 (August 7, 2015): 14–14

100 For a very readable overview of what cognitive neuroscientists have been discovering worldwide about the

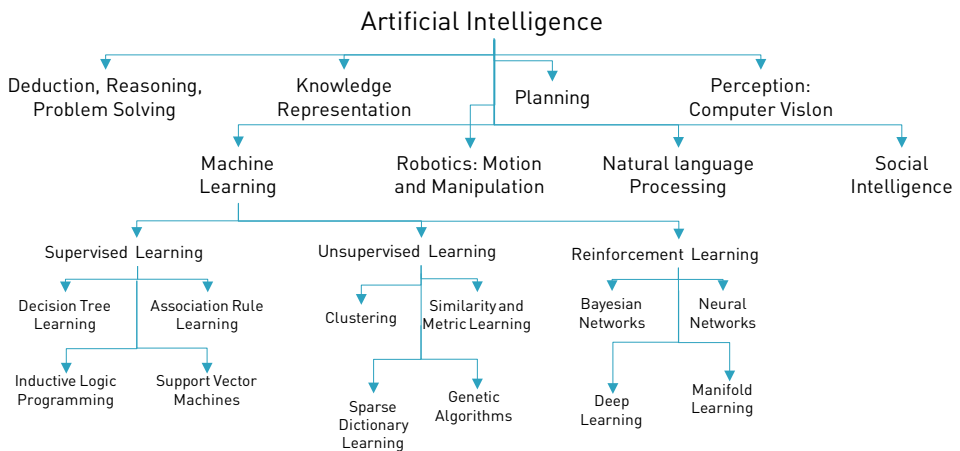


FIGURE 5: AN OVERVIEW OF NOTABLE APPROACHES AND DISCIPLINES IN AI AND MACHINE LEARNING¹⁰¹

By analogy to these components constituting human cognition, what, then, is the library of intelligent sensory-, information-processing-, pattern-matching- and categorization-functions underlying artificial intelligence? As represented in Figure 5, at its highest level, AI problems cluster around several classes. The first of these concerns the problem of correctly **parsing inputs**, which deals with issues of perception; computer vision; natural language processing; and taking appropriate cues from social intelligence; the second of these involves correctly **planning & executing outputs** (‘behavior’), which involves appropriate processes for knowledge representation, prioritization, planning and, in embodied AI systems, robotics (system motion; collision avoidance; manipulators).

2.3.2. Machine Learning: 5 Schools of Thought


Within the distinct subfields of AI, however, there is one field which has been the most responsible for the recent advances in AI system implementations which have driven this third, sustained AI spring: this field revolves around approaches to achieve **machine learning** – helping an AI system learn to identify deep, hidden patterns in existing datasets, or learn to match specific features in data with specific responses or outputs.¹⁰² There are, broadly speaking, **five main ‘schools’ of artificial intelligence machine learning**, taking their main inspiration from different fields of science.¹⁰³

various brain dynamics that ultimately lead to a conscious state, see Stanislas Dehaene, *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts* (New York, New York: Viking, 2014).

101 Based on Nazre and Garg, “A Deep Dive in the Venture Landscape of Artificial Intelligence and Machine Learning.”, slide 3

102 Note that these problem classes frequently overlap: computer vision frequently utilizes machine learning to achieve the pattern-matching of imagery with certain tags. However the key point is that such machine learning approaches are the more broadly applicable learning functions. Computer vision, as an ‘input’ problem, is therefore a subset of the ‘learning problem’ of machine learning.

103 This typology is derived from the work of machine learning expert Pedro Domingos. Cf. Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*, 1 edition (New York:



The first, the aforementioned *connectionists*, take their cue from neuroscience, seeking to study how neurons work and encode knowledge by strengthening synapses between them. As a result, they have developed the neural-net algorithm called backpropagation, which learns from user-provided samples ('correct' input/output pairs). The neural net feeds inputs (such as pictures) through layers of artificial neurons until they emit a final output, then checks whether this output is in accordance with the pre-specified ('correct') output in the training data, a process which is continued until the algorithm becomes able to correctly predict the correct output even for new inputs (in effect demonstrating inductive logic).

Having revived the once-defunct research agenda which originated in the work on perceptrons, connectionists are currently arguably the most famous and successful school, having pioneered the backpropagation learning algorithms which, under the name of 'deep learning', have driven much of the success and applications pioneered by AI giants such as Google, Facebook and Baidu – and which are used for image- and voice recognition, as well as choosing search results and ads. There is an irony in the fact that some of the greatest successes in the quest for non-human intelligence have therefore been driven by attempts which rely on a biomimicry which seeks to approximate, or at least draws inspiration from, the human neural architecture. Nonetheless, in spite of the connectionist success, there are still considerable limits to deep learning systems, as argued by other camps in machine learning. As such, '*Evolutionists*' take their cue from evolutionary biology, and believe they can build AI by emulating (and speeding up) natural selection in digital environments, through genetic algorithms.

Evolutionary programming has had a long and storied history in adaptive design, assembling patents from the generation of non-intuitive designs for devices such as radio receivers, amplifiers and, most recently, 3D printers.¹⁰⁴

Conversely, other machine-learning researchers argue that instead of biology (whether neuroscience or evolution), we should instead work from 'first principles' in computer science and logic. As such, *Bayesians* are a school of thought which is inspired by statistics; Bayesian networks (or 'belief networks') try to encode probability estimates for a large number of different competing hypotheses, with respective belief probabilities updated as new information becomes available.¹⁰⁵ Though capable of being deployed separately,

Basic Books, 2015].; and Pedro Domingos, "The Race for the Master Algorithm Has Begun," *WIRED UK*, 2016, <http://www.wired.co.uk/article/master-algorithm-pedro-domingos>.. For an illuminating interview, see Pedro Domingos, The Knowledge Project: Pedro Domingos on Artificial Intelligence, interview by Shane Parrish, Podcast, August 30, 2016, <http://theknowledgeproject.libsyn.com/pedro-domingos-on-artificial-intelligence>.. See also Cade Metz, "AI's Factions Get Feisty. But Really, They're All on the Same Team," *WIRED*, 2017, <https://www.wired.com/2017/02/ais-factions-get-feisty-really-theyre-team/>..

104 Cf. Domingos, "The Race for the Master Algorithm Has Begun.". For an overview of this field, see John R. Koza, "Survey of Genetic Algorithms and Genetic Programming," in *WESCON/95. Conference record. Microelectronics Communications Technology Producing Quality Products Mobile and Portable Power Emerging Technologies* (IEEE, 1995), 589, <http://ieeexplore.ieee.org/abstract/document/485447/>.; for an account of its history, see Stephanie Forrest and Melanie Mitchell, "Adaptive Computation: The Multidisciplinary Legacy of John H. Holland," *Communications of the ACM* 59, no. 8 (August 2016), <http://web.cecs.pdx.edu/~mm/jhh-cacm.pdf>

105 cf. Judea Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988)

Bayesian networks can also enhance the decision making ability (and increase the transparency) of connectionist neural networks, and as such will likely be incorporated in the designs of self-driving cars and other AI-guided systems – systems for which it is important that the cause of an accident or error can be clearly ascertained.¹⁰⁶

The fourth school, *symbolists*, is closest to the classic strand of knowledge-based AI, still pursuing a general-purpose learning algorithm which can freely combine rules and fill in the gaps in its knowledge. Ultimately, symbolic learning aims at mimicking the thought process of scientists themselves; looking at data to formulate hypotheses, testing these hypotheses against the data to refine them, and deducing new knowledge.

For instance, in 2014 ‘Eve’, an artificially-intelligent ‘robot scientist’ based on symbolic learning and equipped with basic knowledge of molecular biology, made an autonomous discovery, showing that a compound with established anti-cancer properties might also help in the fight against malaria.¹⁰⁷

Finally, the last school in machine learning is *analogisers*. These systems, inspired by human psychology, seek to operate on the basis of analogy – to match new cases with the most similar such situation which has been encountered in the past.

Ultimately, all of these schools have their own strengths, and encounter their own challenges, in their mutual pursuit of a single ‘master algorithm’: the foundation of machines truly able to learn in chaotic, new environments. It is certainly clear that we are still some time away from true or ‘general’ artificial intelligence, and there may yet be hitches or barriers for many of these approaches. And yet, the sheer pace of development over the past few years, the qualitative improvements in the perception-, learning- and decision-capabilities of current systems, and the manner in which the different approaches to AI are increasingly able to catalyze and strengthen each other’s research, all strongly suggest that something has shifted in the field of AI. One way or another, we are now on a highway towards increasingly more capable forms of machine intelligence – with all the danger and opportunity that entails.

106 Cade Metz, “AI Is About to Learn More Like Humans—with a Little Uncertainty,” *WIRED*, 2017, <https://www.wired.com/2017/02/ai-learn-like-humans-little-uncertainty/>.

107 Kevin Williams et al., “Cheaper Faster Drug Development Validated by the Repositioning of Drugs against Neglected Tropical Diseases,” *Journal of The Royal Society Interface* 12, no. 104 (March 6, 2015): 20141289, doi:10.1098/rsif.2014.1289. See also University of Cambridge, “Artificially-Intelligent Robot Scientist ‘Eve’ Could Boost Search for New Drugs,” *University of Cambridge*, February 4, 2015, <http://www.cam.ac.uk/research/news/artificially-intelligent-robot-scientist-eve-could-boost-search-for-new-drugs>; Andy Extance, “Robot Scientist Discovers Potential Malaria Drug,” *Scientific American*, February 2015, <https://www.scientificamerican.com/article/robot-scientist-discovers-potential-malaria-drug/>.



3 AI – TODAY AND TOMORROW

- 3.1 Core Technologies and Applications**
- 3.2 Markets Amidst the Artificial Intelligence ‘Insights Revolution’**
- 3.3 The Timing of AI Development: Future Scenarios**
- 3.4 Disruptors & Wild Cards**
- 3.5 Caveats: Ethical Concerns over AI**
- 3.6 Summing up – 3 Generations of AI**

3. AI – Today and Tomorrow

In just a few years, AI has become far more advanced and 'human-like' than many of us realize, understand – or would like to admit. To take a small selection: in just the past three years, AI systems have proven that they can meet – and exceed – human performance in image recognition,¹⁰⁸ speech transcription¹⁰⁹ and direct translation;¹¹⁰ they have learned how to drive,¹¹¹ identify relevant information in a paragraph to answer a question,¹¹² recognize human faces (even if pictures are blurred) and human emotions,¹¹³ create their own encryption schemes and detect malware¹¹⁴; detect crop diseases,¹¹⁵ write cookbook recipes, sports news articles, movie screenplays, music and published poetry¹¹⁶ – and are able to find their way around the London Underground using a map.¹¹⁷ Do we recognize this? Not always. John McCarthy, who coined the term 'artificial intelligence' back in

108 Allison Linn, "Microsoft Researchers Win ImageNet Computer Vision Challenge," *Next at Microsoft*, December 10, 2015, <https://blogs.microsoft.com/next/2015/12/10/microsoft-researchers-win-imagenet-computer-vision-challenge/>

109 Wayne Xiong et al., "Achieving Human Parity in Conversational Speech Recognition," *arXiv Preprint arXiv:1610.05256*, 2016, <https://arxiv.org/abs/1610.05256>.

110 Davide Castelvecchi, "Deep Learning Boosts Google Translate Tool," *Nature News*, accessed February 25, 2017, doi:10.1038/nature.2016.20696.; for an excellent in-depth account, see also Gideon Lewis-kraus, "The Great A.I. Awakening," *The New York Times*, December 14, 2016, <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>

111 Ross Bryant, "Google's AI Becomes First Non-Human to Qualify as a Driver," *Dezeen*, February 12, 2016,

<https://www.dezeen.com/2016/02/12/google-self-driving-car-artificial-intelligence-system-recognized-as-driver-usa/>

112 Cade Metz, "Google's Hand-Fed AI Now Gives Answers, Not Just Search Results," *WIRED*, 2016, <https://www.wired.com/2016/11/googles-search-engine-can-now-answer-questions-human-help/>

113 LilHay Newman, "AI Can Recognize Your Face Even If You're Pixelated," *WIRED*, 2016, <https://www.wired.com/2016/09/machine-learning-can-identify-pixelated-faces-researchers-show/>; Matt Burgess, "Microsoft's AI Can Detect Your Emotions (but Only If You're Angry)," *WIRED UK*, 2015, <http://www.wired.co.uk/article/microsoft-predict-emotions-artificial-intelligence>

114 Martin Abadi and David G. Andersen, "Learning to Protect Communications with Adversarial Neural Cryptography," 2016, <https://arxiv.org/pdf/1610.06918v1.pdf>; Linda Musthale, "How to Use Deep Learning AI to Detect and Prevent Malware and APTs in Real-Time," *Network World*, March 11, 2016, <http://www.networkworld.com/article/3043202/security/how-to-use-deep-learning-ai-to-detect-and-prevent-malware-and-apt-in-real-time.html>

115 Dyllan Furness, "AI in Agriculture? Algorithms Help Farmers Spot Crop Disease like Experts," *Digital Trends*, October 6, 2016, <http://www.digitaltrends.com/computing/ai-crop-disease/>

116 Alexandra Kleeman, "Cooking with Chef Watson, I.B.M.'s Artificial-Intelligence App," *The New Yorker*, 2016, <http://www.newyorker.com/magazine/2016/11/28/cooking-with-chef-watson-ibms-artificial-intelligence-app>;

Hyacinth Mascarenhas, "Associated Press to Expand Its Sports Coverage by Using AI to Write Minor League Baseball Articles," *International Business Times UK*, July 5, 2016, <http://www.ibtimes.co.uk/associated-press-expand-its-sports-coverage-by-using-ai-write-minor-league-baseball-articles-1568804>;

HAL 90210, "This Is What Happens When an AI-Written Screenplay Is Made into a Film," *The Guardian*, June 10, 2016, sec. Technology, <https://www.theguardian.com/technology/2016/jun/10/artificial-intelligence-screenplay-sunspring-silicon-valley-thomas-middleditch-ai>; Alex Marshall, "From Jingles to Pop Hits, A.I. Is Music to Some Ears," *The New York Times*, January 22, 2017, <https://www.nytimes.com/2017/01/22/arts/music/jukedeck-artificial-intelligence-songwriting.html>;

Brian Merchant, "The Poem That Passed the Turing Test," *Motherboard*, 2015, https://motherboard.vice.com/en_us/article/the-poem-that-passed-the-turing-test; see the original announcement at Zackary Scholl, "Turing Test: Passed, Using Computer-Generated Poetry," *Raspberry Pi AI*, January 24, 2015, <https://rpi.ai.wordpress.com/2015/01/24/turing-test-passed-using-computer-generated-poetry/>

117 Elizabeth Gibney, "Google's AI Reasons Its Way around the London Underground," *Nature News*, 2016, doi:10.1038/nature.2016.20784

1956, quipped that “as soon as it works, no one calls it AI anymore.”¹¹⁸ But still it is there – becoming an integral part of our lives: from personal digital assistants on our mobile phones, to cars routinely outfitted with smart systems that decide when to activate anti-lock brakes, and how to alter the parameters of the fuel injection system. Algorithms guide our plane towards its gate, and determine the price of our flight ticket.

Even something as mundane and underappreciated as email spam filters are now increasingly powered by algorithms which learn from scanning large quantities of email traffic. One key ingredient here is that beyond just learning from all of us collectively, these algorithms are increasingly responsive to our own individual patterns and behaviors. As Google’s John Rae-Grant, a senior product manager for Gmail, states: “One person’s spam might be another person’s coupon”¹¹⁹. Such highly specialized ANI systems work together to inform different platforms online about who you are, and what you might like – and what you might be persuaded to buy. And applications such as recommendation systems and targeted advertisements are just in the consumer world. Sophisticated ANI systems are widely used in sectors and industries including the military, manufacturing, and finance, and medicine. Already, we live amidst an AI revolution – and we hardly even see it.

3.1. Core Technologies and Applications

As has already been demonstrated, the scope for current applications in AI is tremendous. We can distinguish both core technologies, and sectors of applications. Core technologies cover a range of categories, including deep learning & machine learning; natural language processing (both general and speech recognition); computer vision & image recognition; gesture control; virtual personal assistants; smart robots; context aware computing; and speech to speech translation.¹²⁰ AI applications usually fall into the following three categories.¹²¹

In the first place, in terms of **information aggregation**, integration & analysis, AI systems currently include search engines, sentiment analysis (opinion mining), speech and handwriting recognition, spoken language understanding and interface, stock market analysis; health monitoring; news categorization & weather prediction.

In the second place, in terms of practical **tools**, AI systems currently are used for face detection, spam filters, derivatives training, game playing, software testing & automatic

118 Quoted in Moshe Y. Vardi, “Artificial Intelligence: Past and Future,” *Communications of the ACM* 55, no. 1 (January 1, 2012): 5–5, doi:10.1145/2063176.2063177

119 Cade Metz, “Google Says Its AI Catches 99.9 Percent of Gmail Spam,” *WIRED*, July 9, 2015, <http://www.wired.com/2015/07/google-says-ai-catches-99-9-percent-gmail-spam/>.

120 Venture Scanner, “Artificial Intelligence Q1 Update in 15 Visuals,” March 22, 2016, <https://insights.venturescanner.com/2016/03/22/artificial-intelligence-q1-update-in-15-visuals/>

121 These categories, and many of the examples, are derived from Nazre and Garg, “A Deep Dive in the Venture Landscape of Artificial Intelligence and Machine Learning.”, slide 5. See also Bostrom, *Superintelligence*

cyber-vulnerability testing, machine translation, medical diagnosis, hearing aids; mood analysis, brain-machine interfaces (in prosthetics), optical character recognition, recommendation systems, and robotic locomotion.

Finally, AI-enabled **services** include targeted advertising and customer segmentation, classifying DNA sequences, computer-vision object recognition, bioinformatics and chemical analysis, and legal case research¹²² (a field in which it has achieved considerable accuracy: in 2016, a AI developed at University College London was able to predict the judicial decisions of the European Court of Human Rights with 79% accuracy).¹²³

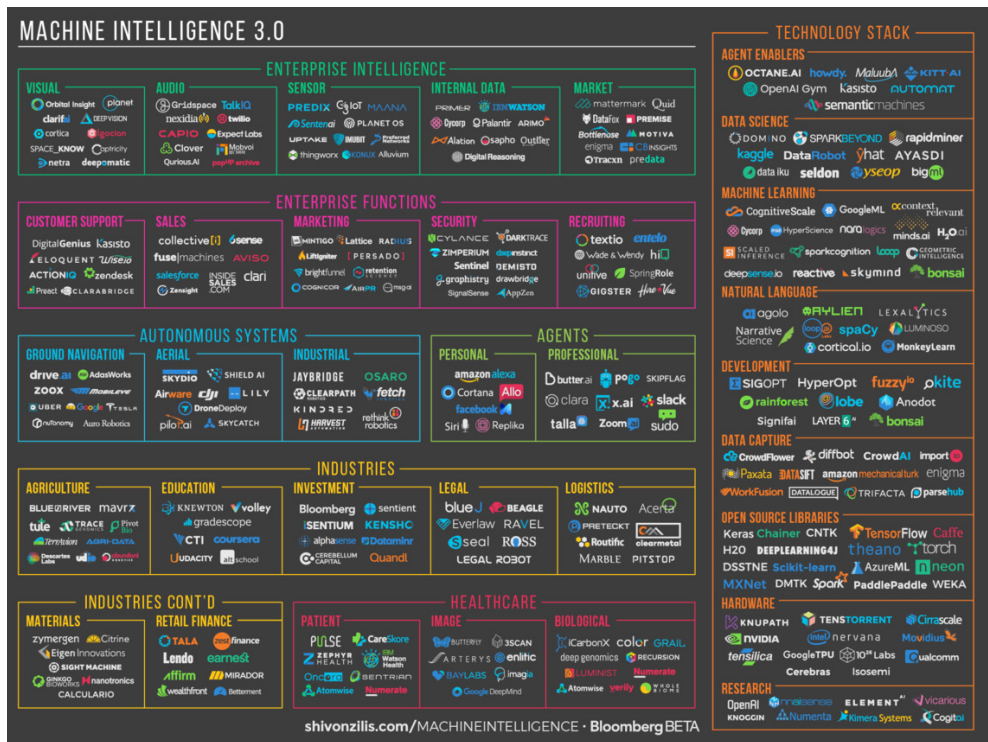


FIGURE 6: AN OVERVIEW OF THE COMPETITIVE MARKET LANDSCAPE FOR MACHINE INTELLIGENCE.¹²⁴

AI also has great potential for transforming government policy; as governments have access to vast quantities of data on diverse sectors, they would be in a good position to deploy machine learning systems to great effect to tailor policy in a range of sectors, from education to medicine, and from effective policing to anticipating reoffense rates for

122 Cf. Cecille de Jesus, "AI Lawyer 'Ross' Has Been Hired By Its First Official Law Firm," *Futurism*, May 11, 2016, <https://futurism.com/artificially-intelligent-lawyer-ross-hired-first-official-law-firm/>.

123 Bex Caygill, "AI Predicts Outcomes of Human Rights Trials," *UCL*, October 2016, <http://www.ucl.ac.uk/news/news-articles/1016/241016-AI-predicts-outcomes-human-rights-trials>

124 Shivon Zilis, "The Current State of Machine Intelligence 3.0," *Shivon Zilis*, 2016, <http://www.shivonzilis.com/machineintelligence>

bail decisions.¹²⁵ Provided that such implementations can be done in a legitimate manner, these programs could increase response time and proactive intervention, while saving massive amounts of scarce public funds.

3.2. Markets amidst the Artificial Intelligence ‘Insights Revolution’

Andrew Ng, former chief scientist at Baidu, has observed that “I have a hard time thinking of an industry we cannot transform with AI.”¹²⁶ Indeed, even a cursory glance will demonstrate that while the situation differs from country to country, on a global level investors have not been blind to the broad utility of AI systems – as a result of which there has been tremendous, sustained investment in the sector,¹²⁷ with a total of more than \$17bn in AI investments between 2009-2014.¹²⁸ As of 2016, about \$2.4billion in private sector funding in AI in 2016 alone.¹²⁹

As a result of investment, there has been a sevenfold explosion in machine-learning patents from 2004-2013; although though most of these patents are not developed by startups, but still held by a few major players such as Microsoft (9%); IBM (6%); Yahoo (4%); Google (3%); and the Nippon Telegraph & Telephone Corp (2%).¹³⁰ Even so, the funding base is diverse, AI research has received more than \$1bn in VC investment since 2010; automotive firms such as Toyota have pledged billions to R&D.¹³¹ Within these markets, a few broad observations apply:

1. In the first place, there has been **a stable trend in AI investment increasing** over time, by double digits, over the last few years.¹³²

125 The Economist, “Machine Learning: Of Prediction and Policy,” August 20, 2016, <http://www.economist.com/news/finance-and-economics/21705329-governments-have-much-gain-applying-algorithms-public-policy?frsc=dgla>.

126 Sarah Zhang, “China’s Artificial-Intelligence Boom,” *The Atlantic*, February 16, 2017, <https://www.theatlantic.com/technology/archive/2017/02/china-artificial-intelligence/516615/>.

127 Dan Barnes, “Investing in Artificial Intelligence - Is It the Right Time to Buy?,” Raconteur, December 15, 2015, <http://raconteur.net/technology/investing-in-artificial-intelligence-is-it-the-right-time-to-buy>. NarrativeScience, “Outlook on Artificial Intelligence in the Enterprise - 2016” (NarrativeScience, 2016), <http://resources.narrativescience.com/h/i/272429469-outlook-on-artificial-intelligence-in-the-enterprise-2016-research-report>. For an overview, see Daniel Faggella, “Valuing the Artificial Intelligence Market, Graphs and Predictions for 2016 and Beyond,” *TechEmergence.com*, March 7, 2016, <http://techemergence.com/valuing-the-artificial-intelligence-market-2016-and-beyond/>.

128 Bank of America Merrill Lynch, “Robot Revolution - Global Robot & AI Primer.”; cf. Gitta Rohling, “Artificial Intelligence: Facts and Forecasts - Boom for Learning Systems” (Siemens, 2014), <http://www.siemens.com/innovation/en/home/pictures-of-the-future/digitalization-and-software/artificial-intelligence-facts-and-forecasts.html>

129 Jason Furman, “Is This Time Different? The Opportunities and Challenges of Artificial Intelligence” (AI Now: The Social and Economic Implications of Artificial Intelligence Technologies in the Near Term, New York, July 7, 2016), https://www.whitehouse.gov/sites/default/files/page/files/20160707_cea_ai_furman.pdf.

130 Andrew Trabulsi, “Future of Artificial Intelligence,” June 2015, http://www.fujitsu.com/us/Images/Panel1_Andrew_Trabulsi.pdf.

131 Mark Holman, “Artificial Intelligence and the Future: Promise or Peril,” June 2, 2016, https://www.atkearney.com/documents/7789801/0/BB_54028_20160212+SVTLS+AI_FINAL.pdf/94492e02-e5c3-4ffb-9e3a-47c9ac8bf93b

132 Rohling, “Artificial Intelligence.” Daniel Faggella, “Artificial Intelligence Industry – An Overview by Segment,” *TechEmergence.com*, July 25, 2016, <http://techemergence.com/artificial-intelligence-industry-an-overview-by-segment/>.

2. In the second place, and perhaps remarkably, an examination of the most well-funded AI startups over the last few years suggests that, at least in the US, **'long-shot' AI platforms and core technologies are receiving more investment than companies focusing on solving specific technology problems in the short run**, suggesting a very bullish and optimistic outlook on the part of venture capitalists.¹³³
3. In the third place, the **IP ecosystem in AI is remarkably diverse**, with machine learning patents running the gamut from text classification; customer analytics; querying methods; medical diagnostics; cybersecurity; data mining; facial recognition.¹³⁴ Nonetheless, it does appear that currently **healthcare, marketing, and finance are amongst the large focus areas** for AI investors.
4. In the fourth place; in spite of broadly shared (if not unanimous) market optimism,¹³⁵ there is also **a large degree of uncertainty amongst investors** about what specific AI applications will actually pay off. The risk is that while AI has the potential to transform business models in almost any field, not all companies are necessarily well positioned (in terms of business model; or in terms of the quantity and quality of data available to them) to become frontrunners. As the economist Robin Hanson has joked: "most firms that think they want advanced AI/ML really just need linear regression on cleaned-up data."¹³⁶ In that context, the term AI is also often overused, leading some to pursue investments that may not have very large pay-offs. As in any market, there will be winners and losers.¹³⁷
5. It appears that along with increasing government concern for the societal impacts of automation,¹³⁸ there is simultaneous **increasing government interest in funding AI innovations in all domains**. In 2015, the US Government's investment in unclassified R&D in AI was approximately \$1.1 billion, and its 2016 'National Artificial Intelligence Research and Development Strategic Plan' called for further expansions of long-term investments in this research agenda.¹³⁹ In

133 Faggella, "Artificial Intelligence Industry – An Overview by Segment."; Daniel Faggella, "Venture Investments in Artificial Intelligence – Trends in 2016 and Beyond – TechEmergence.com," *TechEmergence*, November 29, 2016, <http://techemergence.com/investments-in-artificial-intelligence/>; CB Insights, "Deep Interest In AI: New High In Deals To Artificial Intelligence Startups In Q4 '15," *CB Insights* - Blog, February 4, 2016, <https://www.cbinsights.com/blog/artificial-intelligence-startup-funding-trends/>.

134 Trabulsi, "Future of Artificial Intelligence."

135 Daniel Faggella, "How Investors Feel About Artificial Intelligence – from 29 AI Founders and Executives," *TechEmergence.com*, August 1, 2016, <http://techemergence.com/how-investors-feel-about-artificial-intelligence-from-29-ai-founders-and-executives/>.

136 Robin Hanson, "Good CS Expert Says: Most Firms That Thinks They Want Advanced AI/ML Really Just Need Linear Regression on Cleaned-up Data.," microblog, *Twitter*, (November 28, 2016), <https://twitter.com/robinhanson/status/803302419249725440>. For an expanded version of his argument, see also Robin Hanson, "This AI Boom Will Also Bust," *Overcoming Bias*, 2016, <http://www.overcomingbias.com/2016/12/this-ai-boom-will-also-bust.html>

137 James McCormick, "Predictions 2017: Artificial Intelligence Will Drive The Insights Revolution" (Forrester, 2016), https://go.forrester.com/wp-content/uploads/Forrester_Predictions_2017_-_Artificial_Intelligence_Will_Drive_The_Insights_Revolution.pdf.

138 Cf. US Senate Subcommittee on Space, Science and Competitiveness, "The Dawn of Artificial Intelligence" (2016), <http://www.commerce.senate.gov/public/index.cfm/2016/11/commerce-announces-first-artificial-intelligence-hearing>; for a comparative overview, see also Corinne J. N. Cath et al., "Artificial Intelligence and the 'Good Society': The US, EU, and UK Approach," SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, December 23, 2016), <https://papers.ssrn.com/abstract=2906249>

139 Office of Science and Technology Policy, "The National Artificial Intelligence Research and Development

the wake of the March 2016 AlphaGo victory over Lee Sedol, the South Korean government announced plans to invest 1 trillion won (\$840 million) by 2020 to boost its AI industry.¹⁴⁰ Similarly, the Chinese government has announced plans to develop a 100 billion RMB (\$15 billion) AI market by 2018.¹⁴¹ China is an interesting case, because, like Microsoft and Google, its companies see huge potential in AI. Yet while in terms of fundamental groundbreaking work, they are still lagging behind the US, there has been a massive increase in growth in terms of cited research,¹⁴² and they have proven remarkably adept at rapidly iterating over breakthroughs to develop – and deploy – applications of this technology.¹⁴³ Finally, in March 2017, following estimates by Accenture that AI could add in the region of £654 billion (\$814 billion) to the UK economy by 2035,¹⁴⁴ the British government formulated an ambitious new Digital Strategy which announced £17 million in funding for the sector.¹⁴⁵

6. Finally, in addition to private and government-sponsored research, there are also **non-profit programs**, most notably the OpenAI initiative which, backed by Elon Musk, has pledged \$1bn to develop ‘digital intelligence [...] in the way that is most likely to benefit humanity as a whole’.¹⁴⁶

Estimates for the future trajectory of the industry differ, also depending on the definition of AI utilized. However, virtually all estimates envision stratospheric growth rates as recent breakthroughs in machine learning slowly saturate the market, and are in turn strengthened by even further technological developments: as such, Tractica has estimated that the market for enterprise applications of AI will exceed \$30 billion by 2025, with a particular focus on increasing the accuracy of analysing big data;¹⁴⁷ the Bank of America has predicted that the market could blossom to \$70bn for AI systems (along with \$83bn for robotics) by 2020.¹⁴⁸

Strategic Plan” (National Science and Technology Council, 2016), https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf, pg. 6

140 Lee Chi-Dong, “Gov’t to Invest 1 Tln Won in Artificial Intelligence,” *Yonhap News Agency*, March 17, 2016, <http://english.yonhapnews.co.kr/business/2016/03/17/81/0501000000AEN20160317003751320F.html?bb441f50>

141 Eva Xiao, “The Chinese Government Wants A 100 Billion RMB AI Market By 2018,” *TechNode*, May 27, 2016, <http://technode.com/2016/05/27/chinese-government-wants-100-billion-level-artificial-intelligence-market-2018/>.

142 Niu et al., “Global Research on Artificial Intelligence from 1990–2014.”

143 Zhang, “China’s Artificial-Intelligence Boom.”

144 Mark Purdy and Paul Daugherty, “Why Artificial Intelligence Is the Future of Growth” (Accenture, 2017), https://www.accenture.com/t00010101T000000_w/gb-en/acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth.PDF#zoom=50.

145 Department for Culture, Media & Sport, “UK Digital Strategy” (Gov.uk, March 1, 2017), <https://www.gov.uk/government/publications/uk-digital-strategy>; Department for Culture, Media & Sport, “£17 Million Boost for the UK’s Booming Artificial Intelligence Sector - GOV.UK,” Press Release, Gov.uk, [February 26, 2017], <https://www.gov.uk/government/news/17-million-boost-for-the-uks-booming-artificial-intelligence-sector>

146 Greg Brockman and Ilya Sutskever, “Introducing OpenAI,” *OpenAI*, 2015, <https://openai.com/blog/introducing-openai/>.


147 Tractica, “Early Enterprise Adopters of Artificial Intelligence Will Include the Business Services, Advertising, Media & Entertainment, and Finance & Investment Sectors,” *Tractica*, 2016, <https://www.tractica.com/newsroom/press-releases/early-enterprise-adopters-of-artificial-intelligence-will-include-the-business-services-advertising-media-entertainment-and-finance-investment-sectors/>; Daniel Newman, “Inside Look: The World’s Largest Tech Companies Are Making Massive AI Investments,” *Forbes*, 2017, <http://www.forbes.com/sites/danielnewman/2017/01/17/inside-look-the-worlds-largest-tech-companies-are-making-massive-ai-investments/>.

148 Bank of America Merrill Lynch, “Robot Revolution - Global Robot & AI Primer.”

Finally, research firm Forrester has anticipated that the ‘insights revolution’ enabled by AI will “drive faster business decisions in marketing, ecommerce, product management and other areas of the business by helping close the gap from insights to action”, as a result of which we can expect a greater than 300% increase in investment in AI over 2017 alone.¹⁴⁹ It seems clear that, barring a major wild card (see section 3.4.), the coming years will see a continued bonanza in investment and funding into AI research, likely driving new breakthroughs (both in fundamental design and applications) at a staggering rate, potentially bringing us ever closer to the ‘tipping point’ towards ‘general intelligence’ – or even the prospect of an exponential ‘intelligence explosion’.

3.3. The Timing of AI Development: Future Scenarios

There are essentially two main hypotheses for the medium- and long-term future of AI. The first one could be called a cyclical one. It expects the season-like cycle that has characterized AI so far to continue. This would imply that the current AI Spring (or even Summer)



will be followed once again by another AI winter of unknown duration. The second ‘tipping point’ hypothesis sees us nearing an irreversible tipping point towards an unprecedented (non-human) intelligence explosion. Which of these hypotheses seems more likely?

FIGURE 7: ONE SCENARIO FOR AN INTELLIGENCE EXPLOSION.¹⁵⁰

Conjectures on how quickly AI will develop are “as confident as they are diverse”.¹⁵¹ There is much evidence that the ex-ante predictive accuracy of even experts on anything but some minor and short-term forecasts (and even then – only after observing some common rules of thumb and many caveats¹⁵²) is pretty dismal. Also with respect to specifically AI, an overview article of 95 AI timeline predictions concludes that “expert predictions contradict each other considerably, and are indistinguishable from non-expert predictions and past failed prediction.”¹⁵³

149 Cf. McCormick, “Predictions 2017: Artificial Intelligence Will Drive The Insights Revolution.”
 150 Source image: Tim Urban, “The Artificial Intelligence Revolution: Part 1,” *Wait But Why*, January 22, 2015, <http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html>.
 151 Stuart Armstrong and Kaj Sotala, “How We’re Predicting AI—or Failing to,” in *Beyond Artificial Intelligence* (Springer, 2015), 11–29, http://link.springer.com/chapter/10.1007/978-3-319-09668-1_2.
 152 Philip E. Tetlock and Dan Gardner, *Superforecasting: The Art and Science of Prediction*, First edition (New York: Crown Publishers, 2015)
 153 Armstrong and Sotala, “How We’re Predicting AI—or Failing to.”

Still, there are a number of recent surveys in which several relevant expert communities were asked when human-level machine intelligence (HLMI), defined as “one that can carry out most human professions at least as well as a typical human”,¹⁵⁴ would arrive on the scene. The four groups included participants in a mostly theoretical conference (AI-PT); participants in a few top-level Artificial General Intelligence (AGI) conferences – most of them do primarily technical work; a professional organization of Greek published researchers in the AI field (EEN); and The 100 ‘Top authors in artificial intelligence’ by ‘citation’ in ‘all years’ according to Microsoft Academic Search (TOP100).

	Median	Mean	St. Dev.
PT-AI			
10%	2023	2043	81
50%	2048	2092	166
90%	2080	2247	515
AGI			
10%	2022	2033	60
50%	2040	2073	144
90%	2065	2130	202
EETN			
10%	2020	2033	29
50%	2050	2097	200
90%	2093	2292	675
TOP100			
10%	2024	2034	33
50%	2050	2072	110
90%	2070	2168	342
ALL			
10%	2022	2036	59
50%	2040	2081	153
90%	2075	2183	396

The results, which are reproduced in Figure 8, show a 10% probability of “high-level machine intelligence” (HLMI) – that is, AI that “can carry out most human professions at least as well as a typical human” – by 2022, a 50% probability by 2040, and a 90% probability by 2075. We want to point out that these first two target dates fall well within the time horizons of long-term defense planners – and, arguably even more importantly – of the expected life cycles of some of the big-ticket military items that are currently or will soon be entering our armed forces.

FIGURE 8: POSSIBLE AI DEVELOPMENT TIMELINES¹⁵⁵

154 Vincent C. Müller and Nick Bostrom, “Future Progress in Artificial Intelligence: A Survey of Expert Opinion,” in *Fundamental Issues of Artificial Intelligence* (Springer, 2016), 553–570, http://link.springer.com/chapter/10.1007/978-3-319-26485-1_33.

155 Bostrom, *Superintelligence*. We also note that Ray Kurzweil, Google’s Director of Engineering, has recently moved up his own assessment of when the ‘singularity’ – that point in time when all the advances in technology, particularly in artificial intelligence (AI), will lead to machines that are smarter than human beings – might happen to 2029. Dom Galeon & Christianna Reedy Futurism, “A Google Exec Just Claimed the Singularity Will Happen by 2029,” *ScienceAlert*, 2017, <http://www.sciencealert.com/google-s-director-of-engineering-claims-that-the-singularity-will-happen-by-2029>

3.4. Disruptors & Wild Cards

Several events could significantly accelerate or decelerate progress toward AI, and honesty compels us to admit that we do not know which of them might occur, nor in what order. These include – but are explicitly not limited to:

Increasing difficulty of new breakthroughs. Progress in science depends not just on funding available and the effort put in, but also on how ‘hard’ progress is; some fields see the difficulty of further discovery increase with each successive discovery.¹⁵⁶ AI may prove to be one such field, where recent advances are essentially ‘low-hanging fruit’. Indeed, this ‘intractability’ is sometimes considered to apply to some key subfields of AI already, such as natural language processing¹⁵⁷ or long-term automated planning.¹⁵⁸ There may also be a natural or operational ‘ceiling’ to the intelligence available to neural architectures (even non-biological ones), resulting in diminishing returns of research and an eventual, ‘slowdown’ of our progress towards human-level intelligence.¹⁵⁹

Eventual hardware limitations. On a related note, it is possible that along with conceptual and software limits, we may also reach fundamental physical limits to our hardware: while previous predictions to this effect have not been borne out, it is possible that before long we will reach an end to Moore’s Law, and this will slow progress towards AI. For instance, in 2004, the “serial speed” version of Moore’s Law broke down, and the overall trend in computing power growth was only preserved by making a transition to parallel processors. While this for the moment preserved Moore’s Law, it raised new difficulties for software developers.¹⁶⁰ Overall, the formulation of Moore’s law which has the most economic relevance, *computations per dollar*, continues to hold,¹⁶¹ but it is unclear how much longer this will continue before hard physical limits assert themselves, in a way that cannot cost-effectively be addressed through other innovations.¹⁶²

A breakthrough in cognitive neuroscience. Conversely, while it is currently difficult to map and understand the cognitive algorithms behind human intelligence,¹⁶³ new tools, methods or even concepts might enable cognitive neuroscientists to achieve a quantum

156 Samuel Arbesman, “Quantifying the Ease of Scientific Discovery,” *Scientometrics* 86, no. 2 (February 2011): 245–50, doi:10.1007/s11192-010-0232-6.

157 Ernest Davis, “The Singularity and the State of the Art in Artificial Intelligence,” *Ubiquity* 2014, no. October (November 4, 2014): 1–12, doi:10.1145/2667640

158 Edward Moore Geist, “[Automated] Planning for Tomorrow: Will Artificial Intelligence Get Smarter?,” *Bulletin of the Atomic Scientists* 70, no. 2 (2017): 80–85.

159 Alessio Plebe and Pietro Perconti, “The Slowdown Hypothesis,” in *Singularity Hypotheses*, ed. Amnon H. Eden et al., The Frontiers Collection (Springer Berlin Heidelberg, 2012), 349–65, doi:10.1007/978-3-642-32560-1_17.

160 Samuel H. Fuller and Lynette I. Millett, eds., *The Future of Computing Performance: Game Over or Next Level?* (Washington, D.C: National Academies Press, 2011).

161 See Luke Muehlhauser, “When Will AI Be Created?,” Machine Intelligence Research Institute, May 15, 2013, <https://intelligence.org/2013/05/15/when-will-ai-be-created/>, footnote 9; Ray Kurzweil, *How to Create a Mind: The Secret of Human Thought Revealed*, 1st edition (New York: Viking, 2012), Chapter 10, Footnote 10

162 Chris A. Mack, “Fifty Years of Moore’s Law,” *IEEE Transactions on Semiconductor Manufacturing* 24, no. 2 (2011): 202–7, or Hadi Esmaeilzadeh et al., “Power Limitations and Dark Silicon Challenge the Future of Multicore,” *ACM Transactions on Computer Systems* 30, no. 3 (2012), http://www.cc.gatech.edu/~hadi/doc/paper/2012-tocs-dark_silicon.pdf. For a good overview on the debate on the limits to computation, see Igor L. Markov, “Limits on Fundamental Limits to Computation,” *Nature* 512, no. 7513 (August 14, 2014): 147–54, doi:10.1038/nature13570.

163 Thomas Trappenberg, *Fundamentals of Computational Neuroscience*, 2 edition (Oxford ; New York: Oxford University Press, 2010).

leap in understanding how the human brain gives rise to its own intelligence, in a way which might allow AI scientists to recreate that mechanism in an artificial substrate.¹⁶⁴

Human enhancement. More speculatively, human cognitive enhancement technologies could make researchers (and with them whole academic networks) more effective via cognitive enhancement pharmaceuticals,¹⁶⁵ brain-computer interfaces (the so-called 'neural lace'), or genetic engineering for cognitive traits, speeding up the rate of AGI research, and the exchange of lessons and breakthroughs, at least within communities and companies.

Quantum computing. Having overcome early hurdles and barriers,¹⁶⁶ the next developments in quantum computing nonetheless remain difficult to predict; it is also unclear whether or not breakthroughs in this field could contribute (or conversely, might be necessary) to creating and running advanced AI. However, at present it appears that even if built, a quantum computer might provide dramatic computing speed improvements – but only for specific applications.

A 'Sputnik event' creating large development incentives. The 1957 Soviet launch of Sputnik demonstrated to the world the possibility of space flight – and to the US the possibility of an accurate Soviet nuclear ballistic missile arsenal. The resulting aspirations to explore and claim space, and the fears over a 'missile gap', together sparked a space race between the superpowers, with unprecedented and long-term funding for space projects on both sides. If there is a similar "Sputnik moment"¹⁶⁷ for AI, which demonstrates in a vivid or compelling way that human-level (or smarter-than-human) AI is possible or even imminent, this can drive a sharp race towards the finish line, especially since the winner of this race could reap tremendous economic, scientific, military and geopolitical rewards.¹⁶⁸

Societal collapse – or existential catastrophe. Political, economic, technological, or natural disasters may cause a societal collapse during which progress in AI would be essentially stalled or even reversed.¹⁶⁹ There is a growing body of scientific literature that argues that the risk of extreme, global catastrophes deserve increasing attentions from both researchers and policymakers.¹⁷⁰ This literature distinguishes between

164 For an exploration of resulting scenarios, see Robin Hanson, *The Age of Em: Work, Love, and Life When Robots Rule the Earth* (Oxford, New York: Oxford University Press, 2016).

165 Richard A. Posner, *Catastrophe: Risk and Response* (Oxford; New York: Oxford University Press, 2005); Bostrom and Cirkovic, *Global Catastrophic Risks*.

166 See for instance Nick Bostrom, "Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards," *Journal of Evolution and Technology* 9, no. 1 (2002); Bostrom and Cirkovic, *Global Catastrophic Risks*; T.M. Maher and S.D. Baum, "Adaptation to and Recovery from Global Catastrophe," *Sustainability* 5, no. 4 (2013): 1461–79., and for a policy perspective Sebastian Farquhar et al., "Existential Risk: Diplomacy and Governance" [Global Priorities Project, 2017], <https://www.fhi.ox.ac.uk/wp-content/uploads/Existential-Risks-2017-01-23.pdf>.

167 Bostrom and Cirkovic, *Global Catastrophic Risks*., pg. 2

168 Bostrom, "Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards."

169 Cf. Carl Sagan, "Nuclear War and Climatic Catastrophe: Some Policy Implications," *Foreign Affairs* 62, no. 2 (1983): 257–92, doi:10.2307/20041818.; A.M. Barrett, S.D. Baum, and K.R. Hostetler, "Analyzing and Reducing the Risks of Inadvertent Nuclear War between the United States and Russia.," *Science and Global Security* 21, no. 2 (2013): 106–33.; S.D. Baum, "Confronting the Threat of Nuclear Winter," *Futures* 72 (2015): 69–79.

170 Nick Bostrom and Anders Sandberg, "Cognitive Enhancement: Methods, Ethics, Regulatory Challenges," *Science and Engineering Ethics* 15, no. 3 (September 1, 2009): 311–41, doi:10.1007/s11948-009-9142-5.; Luke Muehlhauser, "When Will AI Be Created?," *Machine Intelligence Research Institute*, May 15, 2013, <https://intelligence>.

'limited' Global Catastrophic Risks (GCRs)—loosely defined as “risk[s] that might have the potential to inflict serious damage to human well-being on a global scale”¹⁷¹—and absolute 'Existential Risks' – “threats that could cause our extinction or destroy the potential of Earth-originating intelligent life”.¹⁷² Such threats, which can include natural risks such as meteorite strikes, volcanic winter, or pandemics as much as anthropogenic catastrophes such as the threat of (inadvertent) nuclear winter¹⁷³ or (with advances in synthetic biology) engineered pandemics. These are extreme tail events which, by their nature, would be hard to predict or prepare for.

Societal distrust and disinclination. Less severely, but still problematic, is the potential for societal disruptions, even if falling short of 'global catastrophic risks', to inhibit technological development. Indeed, some have suggested that ultimately the greatest barrier to the development of AI could be society.¹⁷⁴ As AI systems become ever more powerful, as they automate away more and more human jobs, create pervasive inequality, or if such systems end up used for comprehensive and intrusive government surveillance, some societies may question whether it is wise to create machines more powerful than themselves. Already today, there is a growing international movement against lethal autonomous weapons systems, under the banner of the 'Campaign to Stop Killer Robots'.¹⁷⁵ While the track record of societies willingly relinquishing strategically powerful technologies is hardly promising in this, or perhaps chequered at best, the potential disruption caused by advanced AI systems in economies or in war could be visceral enough for enough people that it strengthens the public case of possible 'AI abolitionists'.

3.5. Caveats: Legal and Ethical Concerns over AI

This last point of societal distrust is relevant as it relates to broader caveats we should bear in mind in the development and deployment of machine intelligence. If the development

[org/2013/05/15/when-will-ai-be-created/](http://www.org/2013/05/15/when-will-ai-be-created/).

171 Eleanor G. Rieffel and Wolfgang H. Polak, *Quantum Computing: A Gentle Introduction*, 1 edition (Cambridge, Mass: The MIT Press, 2011).. See also Russ Juskalian, “Practical Quantum Computers,” *MIT Technology Review* 120, no. 2 (April 3, 2017): 76–81. and Jason Palmer, “Here, There and Everywhere. Quantum Technology Is Beginning to Come into Its Own,” *Economist*, March 9, 2017, <http://www.economist.com/technology-quarterly/2017-03-09/quantum-devices>

172 In the context of AI development, the term 'Sputnik moment' derives from Ben Goertzel's 2011 interview with Hugo de Garis, see Ben Goertzel and Hugo de Garis, “Seeking the Sputnik of AGI - H+ Media,” *H+ Media*, March 30, 2011, <http://hpluismagazine.com/2011/03/30/seeking-the-sputnik-of-agi/>.. It was later incorporated in Ben Goertzel and Joel Pitt, “Nine Ways to Bias Open-Source AGI Toward Friendliness,” *Journal of Evolution and Technology* 22, no. 1 (2012): 116–31.

173 Cf. Muehlhauser, “When Will AI Be Created?,” May 15, 2013., footnote 11

174 David J. Chalmers, “The Singularity: A Philosophical Analysis,” *Journal of Consciousness Studies* 17 (2010): 7–65.; Marcus Hutter, “Can Intelligence Explode?,” arXiv:1202.6177 [Physics], February 28, 2012, <https://arxiv.org/pdf/1202.6177.pdf>, pg. 6

175 Campaign to Stop Killer Robots,” accessed March 21, 2017, <https://www.stopkillerrobots.org/>; Human Rights Watch | 350 Fifth Avenue, 34th Floor | New York, and NY 10118-3299 USA | t 1.212.290.4700, “The Growing International Movement Against Killer Robots,” *Human Rights Watch*, January 5, 2017, <https://www.hrw.org/news/2017/01/05/growing-international-movement-against-killer-robots>.. For an in-depth popular treatment, see Sarah A. Topol, “Killer Robots Are Coming And These People Are Trying To Stop Them,” *BuzzFeed*, accessed March 21, 2017, <https://www.buzzfeed.com/sarahatopol/how-to-save-mankind-from-the-new-breed-of-killer-robots>.

of AI, and its introduction into society, is rushed or mishandled, public concerns over technological unemployment, machine bias, automated surveillance or computational propaganda can and will create critical legitimacy problems, driving public distrust of, and even societal backlash against AI.

Therefore, even if we may still be some time from developing ‘full’ or ‘general’ AI, it is important that we now already think about the legal and ethical implications, and consider measures for the responsible supervision, regulation, and governance of the design and deployment of AI systems – all the while keeping in mind that this is not a zero-sum game, and AI can, on balance and if handled properly, be a beneficial technology with innumerable positive applications.

Thus we might meet the call, by prominent AI researchers, to also commit to the “careful monitoring and deliberation about the implications of AI advances for defense and warfare, including potentially destabilizing developments and deployments.”¹⁷⁶ In doing so, policymakers can take stock of the influential long-and short term research priorities set out by AI experts in the 2015 open letter on “Research Priorities for Robust and Beneficial Artificial Intelligence”;¹⁷⁷ or the principles of guiding ethical research and design, ensuring integrity of personal data and individual access and control, as well as economic and humanitarian issues, as recently set out by the Institute of Electrical and Electronics Engineers (IEEE) in their study on ‘Ethically Aligned Design: a Vision for Prioritizing Human Well Being with Artificial Intelligence and Autonomous Systems.’¹⁷⁸ Perhaps most importantly, we must keep in mind the ‘AI Principles’ agreed to by AI experts at the recent landmark ‘Asilomar Conference on Beneficial AI’ (see Appendix I).¹⁷⁹

3.6. Summing up – The AI Tipping Point

Although we have shown expert and non-expert assessments to differ, our own reading of the literature still leads us to the conclusion that there are good grounds to accept the ‘tipping point hypothesis’. We would point to the following key difference with the previous cyclical downturns:


- **The level and diversity of funding** – the ICT revolution has spawned a number of gigantic new agents that are all putting various aspects of AI at the core of their own plans for the future. It is hard for most of us to wrap our minds around the awesome amount of financial and – especially – intellectual horsepower that

176 2015 Study Panel, “Artificial Intelligence and Life in 2030,” One Hundred Yearly on Artificial Intelligence (Stanford, 2016), https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf.

177 Stuart Russell, Daniel Dewey, and Max Tegmark, “Research Priorities for Robust and Beneficial Artificial Intelligence,” *AI Magazine*, 2015, https://futureoflife.org/data/documents/research_priorities.pdf?x90991

178 IEEE, “Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems,” IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems (Institute of Electrical and Electronics Engineers (IEEE), 2017), http://standards.ieee.org/develop/indconn/ec/ead_v1.pdf.

179 Future of Life Institute, “Asilomar AI Principles,” *Future of Life Institute*, 2017, <https://futureoflife.org/ai-principles/>.



companies like Apple, Amazon, Baidu, Google, Facebook, IBM, Microsoft etc. have at their disposal. At the same time, also more traditional top-spenders on R&D in more traditional industries like automotive (e.g. for driverless cars) or pharmaceutical (e.g. for drug discovery and development) ones are prioritizing AI. Barring some unexpected game-changers, it is hard to see how the momentum that is currently being created might be reversed.

- **The mode of cumulative knowledge building**¹⁸⁰ – for most of the industrial age, technological innovation was stimulated (sometimes even semi-coercively, as in the Manhattan project) by industrial age agents like traditional companies or governments, predominantly within ring-fenced research communities. This mode of cumulative knowledge building still has many strengths, but we see it increasingly being replaced by radically new, much more ‘open’ and significantly faster models of knowledge dissemination and accumulation. Many observers continue to be perplexed on how major AI players like Google or Facebook are just open-sourcing some of their most powerful and creative AI algorithms. But they do – and this is leading to unprecedentedly compressed development cycles. With barriers to access so low, myriad startups across the world explore and mine the various algorithmic treasure troves that the ecosystems of the larger players are putting out there.
- **Big data** – we have (purposefully) stayed away from big data because we primarily wanted to focus on the algorithm-side of this revolution. But the incessantly accelerating deluge of real-time, high-resolution data will – in an energetic *paso doble* with the algorithms – almost inevitably trigger new insights
- **The abundance of visible success**: Finally – and in our view most importantly – as we already pointed out, this is the first time in the few boom- and bust-cycles that AI gone through that the real-life results have been so ubiquitously visible and palpable.

As we already noted, many forms of **ANI** are already with us today. Many authors describe, however, that there is a powerful evolutionary dynamic at work that pushes AI to graduate from this level to the next **AGI** one. We can illustrate this by looking at one of the leading companies behind the current trend: Alphabet Inc, the new parent company of Google and several other companies owned by Google co-founders Larry Page and Sergey Brin. Google has been one of the largest investors in AI and it has made remarkable progress in some of the specific areas of AI through various AI-supported applications. One of its earliest – and in all likelihood most powerful and sustained – thrusts in this direction is in the search space – the initial area that put Google on the map and still provides the bulk of its revenues.¹⁸¹ Google Search’s rankings are now powered by various search

180 Much recent literature is quite enamored with the concept of ‘innovation’, even elevating it as a goal to be pursued for its own sake. In our own thinking, we see improvements less in terms of innovation than in terms of knowledge accumulation.

181 On this amazing story, see John Battelle, *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*, Reprint (New York: Portfolio, 2006); David A. Vise and Mark Malseed, *The Google Story* (New York: Delacorte Press, 2005); Eric Schmidt, *Google: How Google Works*, First edition (New York: Grand Central

algorithms that rely on more than 200 unique signals or “clues” that make it possible to guess what the user might really be looking for. These signals include things like the terms on websites, the freshness of content, your region and the (amended) PageRank algorithm that prioritizes web pages based on how many other pages link to them.¹⁸²

The search algorithm itself is constantly improved (under the hood – also to circumvent attempts by Search Engine Optimization players to ‘game’ the system): Google is said to make about 600 changes to its algorithms each year.¹⁸³ In November 2015, Greg Corrado, a senior research scientist at Google, divulged that a “very large fraction” of the millions of queries a second that people type into the company’s search engine were by then interpreted by an AI system nicknamed RankBrain.¹⁸⁴ RankBrain is a brainchild of the company’s team of machine learning researchers that call themselves *Brain Team*.¹⁸⁵ Corrado stated that within a few months RankBrain had become the the third-most important signal contributing to the result of a search query.¹⁸⁶ Describing the differences between RankBrain and the other technologies in the search engine’s algorithm, he said “The other signals [are] all based on discoveries and insights that people in information retrieval have had, but there’s no learning”.¹⁸⁷ It is interesting to note that, as in other areas of AI, (so far) making the algorithms more human-like in their ability to learn and tweak themselves seems to lead to great improvements in their effectiveness.¹⁸⁸

This search-specific ANI algorithm is therefore now learning exponentially more and more every second about every single one of us that uses Google Search (at last count – almost 75% of us globally)¹⁸⁹ and – accordingly – of all of us as one (or more) group(s). This is turning Google Search – and other similar companies like Baidu, Microsoft, Yandex,

Publishing, 2014); Steven Levy, *In the Plex: How Google Thinks, Works, and Shapes Our Lives*, 1st Simon & Schuster hardcover ed (New York: Simon & Schuster, 2011).

182 For an ‘official’ description, see Google, “Algorithms – Inside Search – Google,” accessed August 21, 2016, <https://www.google.com/insidesearch/howsearchworks/algorithms.html>. For a broader (and more insightful) story about Google Search and its broader implications, see John Battelle, *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture* (Portfolio, 2006).

183 Moz, “Google Algorithm Change History - Moz,” Moz, accessed September 2, 2016, <https://moz.com/google-algorithm-change>

184 Google continues to be extremely secretive in its communication strategy, but what we (think we) know about RankBrain is usefully summed up by the highly respected ‘search’-veteran Danny Sullivan Danny Sullivan, “FAQ: All about the Google RankBrain Algorithm,” *Search Engine Land*, June 23, 2016, <http://searchengineland.com/faq-all-about-the-new-google-rankbrain-algorithm-234440>

185 The company’s head of cloud Diane Greene told attendees at the Oktane tech conference on September 30, 2016 that Google Brain is also a physical entity existing inside Google’s massive collection of data centers. «Visiting the data centers. They are just unbelievable. They are acres and acres and you go through a double door and there’s thousands and thousands more servers. There’s not a lot of people ...And we have something called Google Brain, where we do our machine learning with special processors. And there’s the Google Brain,» she says, describing seeing it behind those double doors. *Business Insider*, “The ‘Google Brain’ Is a Real Thing but Very Few People Have Seen It,” *Business Insider*, September 2, 2016, <http://finance.yahoo.com/news/google-brain-real-thing-very-002720762.html>.

186 Jack Clark, “Google Turning Its Lucrative Web Search Over to AI Machines,” *Bloomberg.com*, October 26, 2015, <http://www.bloomberg.com/news/articles/2015-10-26/google-turning-its-lucrative-web-search-over-to-ai-machines>

187 Matthew Capala, “Machine Learning Just Got More Human with Google’s Rankbrain,” *The Next Web*, September 2, 2016, <http://thenextweb.com/artificial-intelligence/2016/09/02/machine-learning-just-got-more-human-with-googles-rankbrain/>

188 Ibid.

189 netmarketshare.com, “Search Engine Market Share,” *Netmarketshare*, July 2016, <https://www.netmarketshare.com/search-engine-market-share.aspx?qprid=4&qpcustomd=0>.

etc. – into the single largest (fast and accelerating) ‘learner(s)’ about humanity and its environment in all of its aspects – including in the defense and security realm. And that is just ‘search’. There is more – much more. Google is also ‘learning’ – equally quickly – more about humans and our environments from its massive AI-driven efforts in spam-detection for its Gmail product;¹⁹⁰ various traffic-related detectors for its driverless car division;¹⁹¹ traffic-jam related predictions for Google Maps; healthcare-related insights from DeepMind Health,¹⁹² previously unknown real-time knowledge about stratospheric wind currents for Project Loon, etc. And then we are not even talking about the various parts of Alphabet Inc. that are working on more ‘fundamental’ AI problems that would traditionally have been dealt with in academia, like natural language processing (and understanding)¹⁹³ or even neural network tools.¹⁹⁴

If we just focus on X’ (formerly called Google X) driverless cars division, for instance, it is hard to see why the data used to optimize the car’s behavior to avoid various accidents would remain restricted to the data the car’s various sensors currently are collecting. Given the exorbitant costs in human lives, money and time that human driving has exacted, it would stand to reason to augment those ‘lessons’ with any other behavioral or environmental lessons that other forms of Google AI may have learned in different ‘specific’ applications. Since our mobile phones are usually in our cars, and they keep collecting various other forms of information, why would we want to enrich the hard information that the car’s sensors are collecting with what the soft knowledge on us that Google Mail, Google Maps, Google Now, Google Translate, etc. have already learned about us? On the other hand, such informational cross-integrations may run up against public opposition or distrust if not implemented in a well-considered and – secured manner.

When we look at AI expansion from a ‘human’ perspective, we see the Alphabet division leaders deciding to interconnect these algorithms to improve their respective value propositions (and bottom-lines). And we see us, the users, as increasingly benefitting from all of these different personalized services that Google is offering at financially


190 Metz, “Google Says Its AI Catches 99.9 Percent of Gmail Spam.”

191 Jeremy Hsu, “Deep Learning Makes Driverless Cars Better at Spotting Pedestrians,” *IEEE Spectrum: Technology, Engineering, and Science News*, February 9, 2016, <http://spectrum.ieee.org/cars-that-think/transportation/advanced-cars/deep-learning-makes-driverless-cars-better-at-spotting-pedestrians>.

192 For an overview of this work, see IBM Research, “IBM Research: A New Kind of Dementia Treatment,” accessed September 20, 2016, <https://www.research.ibm.com/articles/dementia-treatment-diagnosis.shtml>. Also Alexandra König et al., “Automatic Speech Analysis for the Assessment of Patients with Predementia and Alzheimer’s Disease,” *Alzheimer’s & Dementia: Diagnosis, Assessment & Disease Monitoring* 1, no. 1 (March 2015): 112–24, doi:10.1016/j.dadm.2014.11.012.


193 For example: in May 2016, Google announced (and made open-source) the world’s most advanced syntax-parser for English Slav Petrov, “Announcing SyntaxNet: The World’s Most Accurate Parser Goes Open Source,” *Google Research Blog*, May 12, 2016, <https://research.googleblog.com/2016/05/announcing-syntaxnet-worlds-most.html>; and a mere 3 months later (in August 2016) it announced the parser was now available for 40 (!) languages Chris Alberti, Dave Orr, and Slav Petrov, “Meet Parsey’s Cousins: Syntax for 40 Languages, plus New SyntaxNet Capabilities,” *Google Research Blog*, accessed August 21, 2016, <https://research.googleblog.com/2016/08/meet-parseys-cousins-syntax-for-40.html>.

194 Like TensorFlow, a fully open source software library for machine intelligence that puts wide and deep learning at the disposal of any (relatively software-savvy) programmer or even end-user. TensorFlow, “TensorFlow — an Open Source Software Library for Machine Intelligence,” *TensorFlow*, 2016, <https://www.tensorflow.org/>.



(virtually) no cost. When we look at it from the AI perspective, however, we see how all of these ever more interconnected algorithms are increasingly ‘evolving’ towards more artificial general AI that is starting to know more about us as individuals, as groups and as a species than we do ourselves.

One of the biggest challenges for DSOs, in our view, is to determine what role it will play in this evolution. But before we turn our attention to this question, we feel it necessary to insert a brief chapter on what ‘defense’ actually means in this bigger picture.



4 DEFENSE – YESTERDAY, TODAY AND TOMORROW

- 4.1 Epochs and Technologies
- 4.2 Layers of ‘Defense’
- 4.3 AI in Other Armed Forces Today
- 4.4 Use Cases

4. Defense – Yesterday, Today and Tomorrow

Having identified a number of different ‘layers’ (and generations) of AI, we will now propose to also differentiate between a number of different ‘layers’ (and maybe again also evolutionary mutations) of defense and armed force. HCSS has already emphasized in previous work¹⁹⁵ that many of the terms that we use in a contemporary defense and security context are deeply (and – in our opinion – increasingly counterproductively) embedded in the industrial age.¹⁹⁶ When most of us think about the concept of ‘armed force’ today, we conjure up highly hierarchically organized mobile formations of uniformed soldiers equipped with a wide range of industrial-age physical technologies based (mostly) on steel, engines and firepower (tanks, frigates, jet-fighters, etc.), that are employed by national political leaders to control and secure their territory and to defend and/or advance their national goals through the application of (often lethal) industrial-kinetic¹⁹⁷ violence. Their primary *raison d’être* is thought to be to successfully achieve certain operational mission goals against various opponents. In this prevailing view of defense and armed force, AI is typically seen as a new technology that, very much along the lines of previous technological enhancements like stealth or precision strike, might

195 Stephan De Spiegeleire and Peter Essens, “C2 That! Command and Control over Post-Industrial Armed Forces” (15th International Command and Control Research and Technology Symposium: “The Evolution of C2: Where Have We Been? Where Are We Going?”, Santa Monica CA, 2010).

196 We use the term ‘industrial-age’ fairly loosely here – metals had been used since the Bronze Age, and even firearms go back to at least the Renaissance period and became quite common already in the 18th century. For those interested, see Kenneth Chase, *Firearms: A Global History to 1700* (Cambridge UK; New York NY: Cambridge University Press, 2003); Brenda Buchanan, *Gunpowder, Explosives and the State: A Technological History* (Aldershot England; Burlington VT: Ashgate, 2006). But the argument here is that Industrial warfare saw nation-states creating and equipping large armies and navies (and in the 20th century also air forces) based on mass conscription, rapid transportation (first on railroads, then by sea and air) and unprecedented communication (from telegraph to wireless communications). In terms of physical technology, this era saw the rise of rifled breech-loading infantry weapons capable of massive amounts of fire, high-velocity breech-loading artillery, metal warships, submarines, aircraft, rockets and missiles, etc.

197 Since the 1990s the word ‘kinetic’ has been used (often loosely) by military analysts and practitioners as a particular form of application of military force: in ‘kinetic’ and ‘non-kinetic’ ways. This closely corresponds to what some are now calling ‘right of bang’ vs ‘left of bang’: kinetic force (‘right of bang’) is thought to be the lethal application of (usually physical, explosive and destructive) military force; and non-kinetic force is everything else. The actual roots of the word ‘kinetic’, however, merely refer to something that moves. The word comes from the Greek verb *keinain* (“to move”), which in turn goes back to a proto-Indo-European root **kie-neu-*, the suffixed form of the root **keie-* (“to set in motion”). When looked at more broadly, therefore, kinetic force or kinetic strike can probably best be defined more generally as the application of ‘force’ (the ability to make somebody else do something that s/he may not want to do out of their own volition) through some kinetic vector carrying a payload (the ‘effector’) that is applied to the target. Currently the kinetic capabilities that our armed forces possess combine explosive payloads (firepower) with various typically metal-based vectors with an intent to destroy or kill. In our own work, we prefer to use the term ‘industrial-kinetic’ to differentiate that particular industrial-age type of ‘kinetic’ force from potential new kinetic vectors and payloads (cyber, biometric, behavioral, informational, etc.) to achieve similar or new types of effects (e.g. to incapacitate or detain rather than kill).

augment the ability of our current industrial-age military capability bundle. The same industrial-age actors, doing the same industrial-age things with the same industrial-age – but AI-enhanced – instruments.

While this particular image is by now deeply ingrained in our consciousness, ‘armed forces’ have not always looked like their current industrial-age incarnations.¹⁹⁸ Before the nation-state became the primary actor in the international system (a point in history often traced back by political scientists to the Treaties of Westphalia of 1648), ‘armed force’ was exercised by a far more heterogeneous set of actors than just the nation-states (tribes, clans, religious or ethnic groupings, etc.). And the physical incarnation of this force – the defense capability bundle they used – looked quite differently prior to the industrial revolution than it did afterwards. There are many signs that we once again find us at a new epochal tipping point. Our suggestion in this section is that ‘armed force’ is likely to go through another epochal transformation in which AI’s impact is likely to be far more revolutionary than stealth or precision strike.

This chapter will start by taking a closer look at how technologies – both physical and social – have affected the most fundamental – epochal – shifts in human history. It will then apply some of those insights to defense and armed force, in order to come up with a 4-tiered layering of what defense may be morphing into today.

4.1. Epochs and Technologies

We have just argued that our current way of thinking about defense and armed force is still very much anchored in the industrial age which followed the most recent *epochal* revolution in the history of homo sapiens about two centuries ago. We have all learned in our history books how various *physical technologies* sparked the industrial revolution: water mills, the cotton gin, steam engines, etc. These ‘artificial’ machines did indeed dramatically change our polities, societies, economies, individual employment patterns, demography, and even our cultures. But one of the under-appreciated ‘revolutionary’ aspects of the industrial revolution was the amount of socio-technological innovation that co-evolved around the new physical technological innovations that irrupted into our societies.¹⁹⁹ The new physical technologies (machines) of the industrial revolution were

198 Nor do the etymological roots of this word preordain such a narrow definition. The original root of the word ‘arm’ is thought to derive from the Proto-Indo-European (the hypothetical reconstructed ancestral language of the Indo-European language family whose time scale is much debated, but thought to be about 7,000 years ago, see J. P. Mallory and Douglas Q. Adams, *The Oxford Introduction to Proto Indo European and the Proto Indo European World* (New York: Oxford University Press, 2006).) base **ar-*, meaning ‘to fit, to join together. The notion thus seems to be that ‘arms’ implied ‘that which is fitted together’ (Harper, “Online Etymology Dictionary: Arm.”). This suggests that ‘armed force’ merely represents what we today would call ‘capability packages’ (that what could be fitted together) for the purpose of imposing one’s will on others (‘force’). And the precise instantiation of these capability packages typically reflects the Age in which they are used. In pre-historical times, ‘what was fitted together’ was essentially wood, a few primitive ropes and some stones (for clubs, spears, bows, slings). In the Bronze Age, bronze was added to the mix to yield edged metal weapons; the Iron Age added the much more commonly available iron to the mix – and so on until we reach the current industrial-age ‘armed force’ that we now take as the standard.

199 Richard R. Nelson, “Physical and Social Technologies and Their Evolution,” *Économie Appliquée* 56, no. 3

powerful engines of change in their own right. But what really turned everything upside-down were the social ways in which we humans re-organized ourselves to creatively leverage the opportunities embedded in them.

Both of these drivers (physical and social technologies) were powered by human intelligence. Just as in the (individual – but even more so – social – see also the introduction) ‘intelligence explosion’ that accompanied homo sapiens’ march to the top intelligence played the key role, so too did the social technologies of the industrial age arguably prove more consequential than the physical ones. ‘Division of labor’, for instance, was a social technology that revolutionized first manufacturing, and thence all walks of (human) life. Similarly, the modern ‘firm’ is a social technology that only emerged in this very same period (and that has kept evolving ever since): first in Britain as a (manu) factory; then by the 1920s as the large modern business enterprise.²⁰⁰ The ‘nation-states’ as a social technology co-emerged alongside these other social innovations, as did the accompanying social technologies to manage them, like governments, ministries, civil services, educational systems, armed forces, etc. It may seem counter-intuitive to call these ‘technologies’, but just like their physical counterparts, social technologies had to be ‘invented’ by (groups of) humans and were also constantly tinkered with, refined and adjusted.²⁰¹

Many of these industrial-age social technologies still drive and define our everyday lives. Most of us now work for *governments* or *companies*. *Nation states* still shape our economic, political and social activities, our interests (arguably even our values) and our very identities. Many of us think that all of these are ‘normal’ immutable aspects of the human condition. By doing so we forget that they are essentially little more than transient (human-invented) social technologies²⁰² that have only existed for a few generations and are highly likely to be as subject to disruptive change as the physical – let alone digital – technologies around us.

[2003]: 13–32; Carlota Perez, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages* (Cheltenham UK; Northampton MA USA: Edward Elgar Publishing, 2002); Eric Beinhocker, *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics* (Boston Mass.: Harvard Business School Press, 2006).

200 Judit Kapás, “Industrial Revolutions and the Evolution of the Firm’s Organization: An Historical Perspective,” *Journal of Innovation Economics & Management*, no. 2 (2008): 15–33; Joel Mokyr, *The Gifts of Athena: Historical Origins of the Knowledge Economy* (Princeton, [N.J.]: Princeton University Press, 2002); Chris Freeman and Francisco Louçã, *As Time Goes by: From the Industrial Revolutions to the Information Revolution* (Oxford: Oxford University Press, 2002); Alfred D. Chandler, “How High Technology Industries Transformed Work and Life Worldwide from the 1880s to the 1990s,” *Capitalism and Society* 1, no. 2 (January 7, 2006), doi:10.2202/1932-0213.1004.. And since the Third Industrial (ICT) revolution it keeps evolving to flatter and more decentralized companies that are increasingly organized in semi-autonomous project-based teams.

201 In this context, it may be useful to point out that even the very word ‘technology’ had a very different meaning prior to the industrial revolution, when it referred principally to a field of study concerned with the practical arts Eric Schatzberg, “Technik Comes to America: Changing Meanings of Technology before 1930,” *Technology and Culture* 47, no. 3 (2006): 486–512.. “Oldenziel and Marx in fact argue that it attained the status of “keyword” only in the 1930s, and that before this time, issues that historians now discuss in terms of technology were framed in such terms as useful arts, manufacturing, industry, invention, applied science, and the machine... George Crabb’s Universal Technological Dictionary of 1823 defined technology as “a description of the arts, especially those which are mechanical.””

202 Here too etymology proves revealing: the Greek root *tekhne* («art, skill, craft in work; method, system, an art, a system or method of making or doing) comes from the Proto-Indo-European root *teks, meaning «to weave, fabricate, make». The original sense thus seem to refers to human’s skillfully creative activities.

Many authors claim that we are currently moving into a post-industrial age.²⁰³ And here, the lines between physical, digital and social technologies seem to be blurring. Engineers today are moving from the atom-based (physical) world into the bit-based (digital) world: in the US, the ratio of software to hardware engineers is 53% and growing.²⁰⁴ That means that the current-day equivalents of the physical technologies that shaped the industrial age are no longer purely physical but an increasingly physical-digital hybrid. At the same time, however, we are also – once again – witnessing the emergence of a whole range of new social technologies: the ways in which we as humans try to leverage these new hybrid digital-physical technologies to pursue our needs and desires²⁰⁵. Social networks (Facebook, LinkedIn, Whatsapp, Instagram, Snapchat; but also Airbnb, Avvo, Behance, Doximity, Github, GrabCad, Researchgate, Rizzoma, StackOverflow, Quora, Rallypoint (for the military), Uber, Upwork and many other new – increasingly web-based – software programs and/or apps etc.), for instance, represent a fundamental change in the way in which many of us interact socially and professionally.

In contrast to the industrial age, those new social technological changes have not just been triggered by new physical technologies. This time around, it is developments in both

203 The seminal work on this remains the 1976 classic Daniel Bell, *The Coming of Post-Industrial Society: As Venture in Social Forecasting*, Peregrine Books (Harmondsworth: Penguin Books, 1976), which was popularized in Alvin Toffler, *The Third Wave, A Bantam Book* (New York: Bantam Books, 1990) and - applied to the defense realm - in Alvin Toffler and Heidi Toffler, *War and Anti-War Survival at the Dawn of the 21st Century* (New York: Little Brown, 1993). Another important milestone in this discussion was Manuel Castells' work on the 'network society' (Manuel Castells, *The Rise of the Network Society*, 2nd ed., with a new pref, The Information Age : Economy, Society, and Culture, v. 1 [Chichester, West Sussex ; Malden, MA: Wiley-Blackwell, 2010]. Overall, we continue to prefer to characterize what is happening around us now as a 'post-industrial' revolution rather than (merely) a 'fourth industrial' one (or Industry 4.0), even though the 'fourth industrial revolution' meme has been gaining quite a bit of traction since it became the main theme of the 2016 World Economic Forum in Davos. See Klaus Schwab, "*The Fourth Industrial Revolution*" (World Economic Forum Geneva, 2016), <https://foreignaffairs.org/articles/2015-12-12/fourth-industrial-revolution>; Klaus Schwab, *The Fourth Industrial Revolution* (World Economic Forum, 2016). Here is how Klaus Schwab describes 'his' 4 industrial revolutions: "The First Industrial Revolution used water and steam power to mechanize production. The Second used electric power to create mass production. The Third used electronics and information technology to automate production. Now a Fourth Industrial Revolution is building on the Third, the digital revolution that has been occurring since the middle of the last century. It is characterized by a fusion of technologies that is blurring the lines between the physical, digital, and biological spheres." Different experts and scholars use different periodizations for these 'industrial revolutions' (for another one, see Perez, *Technological Revolutions and Financial Capital*). We agree that what is going on now is a genuine revolution, but our reasons for not lumping this new revolution together with the first two (or even with the third one - which we see as far from completed but already carrying in itself the germs of even more fundamental changes that the fourth is likely to bring to fruition) are twofold. First of all, we see the main agents of change behind the more general, canonical industrial revolution as lying not so much in the physical technologies that triggered it (steam, coal, iron production) but in the social ones ('the 'modern' enterprise; 'modern' government; the 'modern' (Weberian) state; 'division of labor', etc.). So too do we surmise that the new 'social' technologies (between people and - to a smaller extent - machines) that Schwab's third revolution is starting to engender, and the increasingly more AI (and not machine-)centric ones that his fourth one is likely to spawn are likely to be dramatically different from the first two. Secondly, we feel that whereas the main organizing principle behind the 'industrial age' was the 'line', the post-industrial age is much more likely to become the age of the 'network' or even the 'ecosystem'. Taking these two arguments together, we feel that calling what is going on now just a fourth installment of a recurring pattern underestimates its far more disruptive nature. We do want to point out, however, that Schwab too sees this 'fourth' installment as a quite different one: "When compared with previous industrial revolutions, the Fourth is evolving at an exponential rather than a linear pace. Moreover, it is disrupting almost every industry in every country. And the breadth and depth of these changes herald the transformation of entire systems of production, management, and governance."

204 Another interesting aspect here is that the typical desired attributes of these engineers is also changing, in the sense that creativity is becoming a more important part in the mix. See also Adam Grant, *Originals: How Non-Conformists Move the World* (Penguin, 2016).

205 McKinsey Global Institute et al., *The Social Economy: Unlocking Value and Productivity Through Social Technologies* (McKinsey Global Institute, 2012).

hardware and software – primarily, but not only in the ICT-sector – that are the key drivers behind emerging new social technologies. But if we still try to differentiate between – on the one hand – what scientists are developing in terms of hard- and software and – on the other hand – how we as members of society are starting to use these scientific-technological developments to reorganize the way in which we create ‘value’, we see that the interaction between these two elements is leading to truly disruptive change – just like it did during the industrial revolution.

In the private sector, we see that this change is leading to some fundamental changes of the main interacting agents *and* the way in which they interact. The (industrial-age) ‘company’ is being reimagined into a set of new social technologies. What some have called the ‘platform’ revolution, for instance, see the industrial-age firm morph into a new business model that uses technology to connect people, organizations, and resources in an interactive ecosystem in which unprecedented amounts of value can be created and exchanged²⁰⁶

This has already occurred in the telecom and media sectors. We are witnessing it in the retail, mobility and hospitality sectors. Is it prudent to assume that ‘defense’ will prove immune against these trends?²⁰⁷ Or should we start anticipating that the ongoing technological acceleration – including, maybe even relatively soon, an upcoming new intelligence explosion – is likely to engender fundamental changes in both physical and – maybe even more importantly – social ‘defense’ technologies? Changes that may not merely change the capability bundle our armed forces currently use to execute their operations, but that may end up changing the very essence of who they are, what they do, and how they do it. In the next section, we will explore what these changes might mean for defense and security organizations.

4.2. Layers of ‘Defense’

Virtually all of the literature on the nexus between AI and defense that we have perused in preparing this think-piece focuses on one particular incarnation of ‘defense’ and ‘armed force’: one that is primarily based on industrial-age physical technologies (violence) as

206 Geoffrey Parker, Marshall Van Alstyne, and Sangeet Paul Choudary, *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*, 2016, <https://www.overdrive.com/search?q=747D962F-32B3-4F05-9126-A09E659AB96C>

207 In their ‘Platform Revolution’, Parker et al. (2016) point to four attributes that make an industry ready for ‘platform’ disruption: being information-intensive; having non-scalable gate-keepers; being highly fragmented (I) and being characterized by extreme information asymmetries Ibid.. Many might argue that defense is different: that it is physical rather than informational; that the ‘gate’ is effectively kept by the nation states’ pre-digital monopoly on legitimate violence; that the defense industry is oligo-polistic and -psonistic and that existing information asymmetries (supported by ‘secrecy’) are desirable and sustainable. Our counter-argument would be that this is an industrial-age interpretation of what defense really is; a mental association that is quite analogous to those in the mobility sector who thought ‘Detroit’ was safe; only to be upended by Silicon Valley. We see defense as increasingly becoming informational; with - poorly scalable - national armed forces as gate-keepers that have an increasingly hard time in ‘holding’ the gate; with a plethora of currently poorly connected potential (security and defense) value providers; and with an increasingly level informational playing field.

the main sources of actionable power. Most of these writings accept that we have now crossed a 'tipping point' in AI and look – logically consistently within that paradigm – for ways to embed this new '(fourth) industrial-age' technology into – essentially – the same types of 'armed force' executing the same types of 'operations' with very similar, but now AI-enhanced 'weapons' and concepts of operations. The main argument of this think piece is that this prevailing logic – the paradigm within which we think about security, defense and armed force – may be increasingly incongruous with the disruptive transformations that are afoot.

We have already referred in the beginning of this report to a temporally distant era and form of 'defense' in which the relatively vulnerable and seemingly weak *homo sapiens* managed to climb to the top of the food chain over many physically far more potent adversaries. It did so, by the looks of it, much more through a combination of cognitive and social technologies than through purely physical ones. Despite our species' historical 'defense' experience over these past few millennia (or even centuries²⁰⁸), we still tend to think that the actor with stronger physical vigor, with superior (industrial) firepower, with more potent protection against that firepower, etc. is likely to gain the upper hand in the contests of will of the future. It is based on that extrapolation that we pour ever more money into those attributes of power for defensive and/or offensive purposes. In this paradigm AI is seen to augment our current weapon systems by giving them ever more physically potent and destructive – and increasingly autonomous – capabilities.

Our predilection for (industrial-age) physical forms of 'defense' is even visible in the ways in which we try to imagine the future of our species' defense efforts in science fiction²⁰⁹. These authors often portray curiously anthropomorphic alien creatures organized along mostly human social technological lines (with governments, separate 'warrior' castes, etc.) moving around the universe in suspiciously – some might say comically – human-like steel vessels ('alien spaceships') and 'shooting' at each other or at us with suspiciously recognizable weapons like missiles or – in the more imaginative cases – photon torpedoes. In projecting such futures, we essentially deny our very own basal human evolutionary trajectory, in which none of these attributes proved either dominant or decisive. Are we cognitively doomed to only 'remember' the future instead of creatively dream it up?²¹⁰ Is it conceivable that we are 'stuck' in a particular mental frame that prevents us from thinking about truly different ways of achieving our aspirational defense and security goals? Is it not more likely, for instance, that (benevolent or malevolent or

208 Most analysts would probably also agree that the West's 'victory' over the Soviet Union had less to do with its 'brute-force' physical military superiority (where the Soviet Union may even have had superior capabilities in quite a few areas) than with its vibrant economy and society that was able to navigate the shift to a technologically far more advanced service economy far more elegantly than the much more industrial-age Soviet Union.

209 For some recent exemplars of science fiction writers imagining the future of warfare, see NATO Allied Command Transformation, *Visions of Warfare 2036*, ed. Trina Marie Phillips and August Cole (Norfolk VA: NATO Allied Command Transformation, 2017); United States Marine Corps, *Science Fiction Futures; Marine Corps Security Environment Forecast: Futures 2030-2045* (Quantico, VA: United States Marine Corps, 2016), http://www.mcwl.marines.mil/Portals/34/Documents/FuturesAssessment/Marine%20Corps%20Science%20Fiction%20Futures%202016_12_9.pdf?ver=2016-12-09-105855-733.

210 See Stephan De De Spiegeleire and Tim Sweijjs, *Volatility and Friction in the Age of Disintermediation. HCSS StratMon 2016-2017 Yearly Report, 2017*, for how humans primarily 'remember' the future

even neutral) alien species would have myriad other ways to achieve their 'policy' goals than by just 'shooting' at us the way we have historically done at one another?

Staying closer to the current day and age – could it be that the current Russian leadership's use of (mostly non-physical) behavioral influencing technologies (primarily targeted at their own population, but also at 'vulnerable' segments in our midst) heralds a more novel – and more (also cost-) effective – approach to defense planning than the West, in all of its (undisputed) technological superiority, has been able to muster? And are we so satisfied with our own attempts over these past few decades to use (increasingly expensive) industrial-kinetic military capabilities to 'stabilize' various parts of the world where we thought our interests and/or values jeopardized?

We have suggested in previous HCSS work that 'defense' and 'armed force' might very well start looking very differently again from what we have grown accustomed to over the past few 'modern' centuries.²¹¹ If the security returns on our industrial-age defense and security investments and efforts seem to be diminishing and if 'power' is becoming increasingly dematerialized, does it still make sense to focus so overwhelmingly on robotic forms of actualizing 'AI-power' to achieve our defense and security aims? Will intelligent algorithms that are as intelligent as or even more intelligent than humans still need industrial-age machines to make us behave in ways that their masters (or increasingly maybe even they themselves) consider optimal from the point of view of their own values?²¹²

Our own intuition is that the advent of AI (and its co-emerging technologies) requires us to look more broadly, probably quite differently, but above all more creatively – at 'defense' and 'armed force' and at the role AI is likely to play in those. Sheer cognitive (and social) intelligence has proved to be far more evolutionarily powerful than brute human or human-invented physical force. As we find ourselves on the cusp of what is likely to become a new intelligence explosion, does it really still make sense to restrict ourselves to presentist or recentist views of what defense and/or armed force really represent?

With this health warning in mind, this study proposes the following four-tiered classification of 'defense' as a social technology:

- 'Defense' as military operators ('warfighters' in the parlance of some nations) – e.g. the Netherlands Defense Force (NDF);
- Defense as an organization that supports these operators but also interacts with its counterparts – e.g. the Netherlands Defense Organization (NDO);
- Defense as a player in an increasingly more whole-of-government security-oriented approach – e.g. the Netherlands Defense and Security Organizations (NDSO); and

211 Stephan De Spiegeleire and Peter Essens, "C2 That! Command and Control over Post-Industrial Armed Forces," in *The Evolution of C2* (15th International Command and Control Research and Technology Symposium, Santa Monica CA, 2010).

212 On the issue of 'values' for/of AI, see Bostrom, *Superintelligence*.

- Defense as the (potential) catalyst of a broader defense and security ecosystem of sensors and effectors – e.g. the Netherlands Defense and Security Ecosystem (NDSE)



FIGURE 9: 4 LAYERS OF DEFENSE

Most of our defense efforts are currently focused on the heart of this bull’s eye: our armed forces, and especially their operational (warfighting) ‘edge’, which is widely thought to embody the core competency of ‘defense’. The various supporting elements of the armed forces and the broader organization around them are primarily seen as supporting this hard core. The third, ‘whole of government’ layer includes all elements of national government that bear some responsibility for national security writ large. This layer tends to remain quite stovepiped in most countries, and therefore relatively embryonic as a cohesive layer – although we can definitely discern a notable (if glacially slow) trend in that direction. In the Netherlands, for instance, this layer has manifested itself in efforts like the ‘3D’ approach around expeditionary operations; on topics like water security, cybersecurity, counterterrorism; or in the whole-of-government efforts around national risk assessment.

The fourth, defense and security ‘ecosystem’ layer remains mostly a concept, even if it is gaining significant traction in the Netherlands.²¹³ This layer contains all actors that have some direct or indirect affinity with the Netherlands and its security and that might be able to be ‘mobilized’ in any defense or security effort. In relatively stable and secure security environments, like the one the Netherlands has by and large enjoyed since the end of the Cold War, this ecosystem remains mostly latent and therefore invisible. But when the security environment becomes more threatening or threatened, the elements of that

213 The defense and Security Ecosystem was the central theme of the major 2017 ‘Future Force Conference’, organized by the Dutch Chief of defense, which brought together some 1200 representatives of the NDSE to think about how to turn this idea into an actionable reality. Ministry of Defence, The Netherlands, The Hague Centre for Strategic Studies, and Future Force Conference 2017, “From Partnerships to Ecosystems: Combining Our Efforts for a More Secure World,” *Future Force Conference*, accessed January 29, 2017, <https://futureforceconference.com/theme/>.

ecosystem become more perceptible as they start ‘rallying around the flag’ domestically or start expressing sympathy for the country from abroad – whether in words or in deeds. These different layers of ‘defense’ of course interact with each other – both within a country and across (friendly and not-so-friendly) countries.

All of this is quite comparable to the business world, where we also have different layers that interact with each other. These interactions have come to be known by some widely used acronyms like B2B (business-to-business) or B2G (business-to-government) or B2C (business-to-customer). Businesses interact with other businesses (B2B) throughout their value chain: the companies they buy things (or – increasingly – services) from, the companies they sell things (or – increasingly – services) to, etc. Many of them also interact with governments (B2G), which continue to represent 40-50% of our GDP in developed countries: they lobby them, they are subject to their regulatory authority, but they also ‘sell’ to them. And finally, and probably most importantly, businesses interact with their end-customers (B2C): they monitor their consumption preferences (increasingly also through AI – see Amazon, Google, etc.), they bombard them with behavioral influence strikes (advertisements for their products), they sell to them as customers, etc.




FIGURE 10: LAYERS OF INTERACTION IN THE BUSINESS WORLD

Recently, IBM’s CEO Ginni Rometty has added a new acronym to this group: B2I – business-to-individual. This goes beyond the more narrowly transactional B2C interaction between a business and an individual. Rometty argues that businesses increasingly have to realize that any individual, beyond being a mere transactional ‘customer’ of their products, is also somebody the company should (and increasingly can) understand and engage with as a broader individual.

She sees IBM making a new dramatic pivot toward what IBM now calls cognitive computing – they prefer calling these new technologies ‘cognitive’ instead of ‘artificial’²¹⁴ –

214 Cognitive computing refers to systems that learn at scale, reason with purpose and interact with humans naturally. Rather than being explicitly programmed, they learn and reason from their interactions with us and from their experiences with their environment.” Kelly, “Computing, Cognition and the Future of Knowing.”

in which a much more 'intimate' relationship between machines/algorithms and humans as full-fledged individuals creates new business opportunities. This line of reasoning comes from the mouth of one company's (IBM) CEO, but we suggest that this thinking is increasingly permeating the thinking and acting of many leading-edge companies, most of which tend to be platform businesses²¹⁵ increasingly with AI at their very cores (including Amazon, Apple, Baidu, Facebook, Google, Microsoft, etc).

The argument that we will be developing here is analogous to what Ginny Rometti is now arguing for the business world, but then applied to the defense and security realm. For a long time – throughout the industrial age really – 'defense' has been about A2A: 'our' armed force against 'their' armed force. One 'order of battle' against another (or multiple other) one(s). We now arguably see a trend towards defense (and security) increasingly becoming about one (or a few allied) 'whole-of-government(s)' (or even – societies) against another (or multiple other) 'whole-of-government(s)' (or even whole-of-society(/ies)): D2G and/or D2S. The 'comprehensive approach' (and/or hybrid warfare) is a good example of this trend in various countries. Recent problems with the public acceptance of the value proposition of our armed forces have also triggered a renewed focus on D2P (politics) and D2C (citizen).

But is it possible that, just as business is becoming more and more about B2I, defense might be becoming more about D2I? That we can bypass many of the institutional middle men,²¹⁶ most purely transactional relations (the citizen as a 'customer' of defense and security) in order to drill down to the actual (complex) individual human being with a unique utility function that polities' primary role would be to foster (human flourishing) and that should become the primary target of our purposive actions?

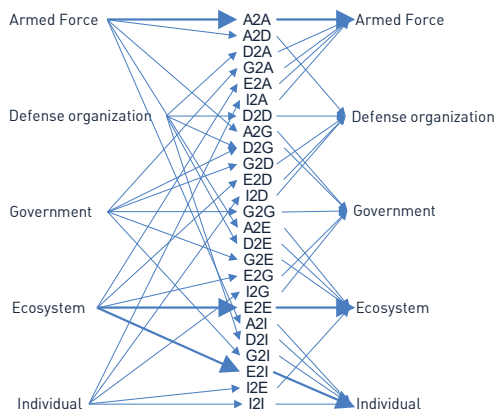


FIGURE 11: DEFENSE AND SECURITY ECOSYSTEM

215 Parker, Van Alstyne, and Choudary, *Platform Revolution*.

216 On the trend towards disintermediation, see our 2016-2017 HCSS Strategic Monitor. De Spiegeleire and Sweijls, *Volatility and Friction in the Age of Disintermediation*. HCSS StratMon 2016-2017 Yearly Report.

4.2.1. Armed Force: AI and A2A

Most of our current discussions about 'defense' (and – consequently – also about the role of AI in defense) are focused on this first – in the current military mindset 'core' – layer of Figure 11. From this vantage point 'defense' is seen as an assembly of individuals that are given a set of of extraordinary rights (also in terms of the tools they are 'allowed' to use) and responsibilities in order to defend our societies against other 'militaries': A2A. The focus in this contest between (or cooperation amongst) armed forces is on 'operations' and 'operators' (in US parlance – 'warfighters').

They embody the sharp, bleeding edge of a broader bureaucratic institutional setup that was created to support it. In line with some of the influential thinking in the business management literature about how companies should just focus on their 'core competencies',²¹⁷ the view here is that 'fighting (and winning) wars' is the core competence of our armed forces – under the political control and strategic guidance of their political leaders. In a nicely linear fashion, this view delineates our militaries as implementation agencies to which the conduct of politically mandated and (hopefully also legally) legitimated military operations is insourced. Consequently, their attention should be squarely and solely focused on the (order of battle of the) 'enemy' and his capability bundle that has to be defeated (preferably decisively) by 'our' order of battle with our capability bundle.

Western defense analytical and political thinking has over the past few decades focused on how to optimally empower the edge of the defense effort under the catchphrase "Power to the Edge"²¹⁸. *In the operational sense*, this implied relying on ever more sophisticated information technologies to better link sensors and effectors ('network-enabled capabilities' in NATO parlance). *In the political sense*, the implication was that the frontline operators had to be as ring-fenced as possible from the various cutbacks that were being inflicted on our defense organizations. Not inconsequentially, this focus on the operators as the leading edge also jibed nicely with the increasingly dominant ruling market-ideology that was supposed to (rightfully – in the eyes of these authors) replace the suffocating bureaucratic administrative-control approach to defense planning²¹⁹.

In this new philosophy, the the real 'customer' that should be at the heart of a more market-driven approach to defense and security was not the political-military decisionmakers in our ministries, but the operator who was supposed to better know what 'the market' required and who had to be better supported by a leaner and more efficient public administration based on better (and more transparent) market incentives.

217 C. K. Prahalad and Gary Hamel, *The Core Competence of the Corporation* (Harvard Business Review, 2001).

218 David S. Alberts and Richard E. Hayes, *Power to the Edge: Command, Control in the Information Age, Information Age Transformation Series* (Washington, DC: CCRP Publication Series, 2003).

219 For an analogous view on the suffocating role of bureaucracy in the private sector, see Gary Hamel and Michele Zanini, "Excess Management Is Costing the U.S. \$3 Trillion Per Year," *Harvard Business Review*, September 5, 2016, <https://hbr.org/2016/09/excess-management-is-costing-the-us-3-trillion-per-year>.


Table (after Harlye 2005)	Public Administration	New Public Management	Design Thinking
Population	Homogeneous	Atomizer	Complex
Needs/problems	Straightforward, defined by 'professionals' / 'experts'	Wants/express by 'the market'	Complex, volatile, prone to risk
Strategy	State- and producercenter	Market - and consumer centered	Service- and consumer centred
Governance through actors	Hierarchies, public servants	Markets, purchasers and providers	Networks and partnership
Key concepts	Public goods	Public choice	Public value
Role of policy markers	Commanders	Announcers / commissioners	Leaders and interpreters
Role of public managers	Clerks and martyrs	Efficiency and market maximizers	Explorers
Role of population	Clients	Customers	Co-producers

FIGURE 12: SHIFTS IN THE DOMINANT MODELS OF THINKING ABOUT PURPOSIVE PUBLIC ACTION

Based on this line of thinking, which has remained dominant throughout most of the industrial age to this date, new technological breakthroughs should be primarily focused on the operational exigencies of military operations. Given the still predominantly industrial-kinetic nature of those operations, this means that the focus should be on our operators and their weapon systems, which are still seen as the most powerful effectors. Based on this premise the logical implication is that AI should be used to improve the (mostly) kinetic capabilities that 'win wars': target detection and acquisition; autonomous weapon systems; planning and support tools, etc.

4.2.2. Defense Organizations: AI and A2D (/D2D)

The second layer of the current defense model is essentially the (bureaucratic) derivative of the first one. The 'sharp edge' (the armed 'operators' or 'warfighters') can only be effective if it is supported by a modern, competent governmental bureaucracy. Even if the 'teeth' are ultimately more important, so goes this line of reasoning, they can only exist by the grace of an efficiently organized supportive 'tail' in turn supported by an efficient Weberian bureaucracy. The focus therefore shifts from the '(operational) edge' to the broader 'defense enterprise'. This approach is surprisingly popular and deeply embedded in our current way of thinking, planning and doing. Part of this is understandable given the large amount of people who are involved in the 'defense enterprise'. Also many other public and private actors see this 'layer' as the interface with (and the gatekeeper of) the 'operators' who have to be protected against, buffered against the 'impurities' of the outside world. From this perspective, it is this more bureaucratic layer that is the real customer. They have the best interest of the operators in mind, but are (presumably) better equipped to deal with all of the other commercial, political, ideological etc. vectors that try to influence them.




It can be expected that this second layer of defense will also start leveraging ever more elements of AI in its everyday practices and procedures. Many of these have become of such byzantine complexity, that various more rule-based software tools have already been introduced over the past few decades – to decidedly mixed effect even within the confines of this layer – to help manage them. Examples here include the introduction of various 100s of legacy software systems that have been rolled out over the past few decades and have recently started being integrated into all-encompassing enterprise resource planning (ERP) systems like SAS. Most of these (often proprietary) systems were still based on the pre-cognitive programming model of human software: engineers writing millions of lines of code that often proved buggy, were not particularly secure and then had to be iteratively ‘fixed’ at great (financial and non-financial) expense.

As more efficient learning algorithms – first still ANI-based, but presumably also increasingly AGI-based – gradually start to taking over from programmed algorithms, it stands to reason that the D2D and D2G layers will represent great opportunities for AI applications. Some low-hanging fruit ANI applications might include recruitment; acquisition; budget management, human resource management, knowledge management, and so on. As in other non-defense-related walks of life, the more large-scale roll-out of such algorithms is likely to have significant effects on the defense enterprise workforce.

4.2.3. Defense and Security Organizations: D2G (/G2G)

The third defense layer extends beyond our defense organizations proper. It includes the entire public sector’s purposive effort to safeguard our societies’ security. A2A experiences over the past two ‘expeditionary’ decades have demonstrated fairly persuasively that it has now become virtually impossible – if ever it even was possible – to achieve one’s defense and security goals through exclusively or even predominantly (industrial-kinetic) military means. To their credit, our armed forces have been amongst the most vocal advocates of this humbling recognition. This has led to a glacially slow but still discernable trend towards more ‘comprehensive’ approaches to defense and security solutions, in which many of the axioms of the previous age had to be – and still are being – rethought.

Our governments (including our defense and security organizations) are still based on various linear industrial-age social technologies that proved extremely useful when the main challenge was to ‘mass’-produce various physical products and to enforce law and order in a (mostly) top-down way. Various digital technologies are exerting increasing pressure on these linear structures and mindsets. If (and it is an ‘if’) it was possible to by and large divorce economic planning from defense planning, for instance, in the past; then today it increasingly is not. This requires our governments to re-invent themselves to reflect the requisite variety of the national and international world that they are trying to affect.



This third layer has two important sub-layers: a more *bureaucratic* one, where the different agencies within our governments have to arrive at better coordinated purposive actions; and a more *political* layer, where the various diverging political interests in our societies also have to find ways to better align themselves in ways that enhance ‘public value’ – also in the defense and security realm. We submit that the potential of AI to intelligently rationalize and align the various inputs, throughputs and outputs of these various processes is enormous. The amount of (increasingly electronic) paperwork that our governments produce is mind-boggling. Most of this is produced in a stovepiped way, with certain government departments producing policy that may be in direct conflict with policy being produced in another government department – without either of them even knowing about the other. AI’s deep learning algorithms offer unprecedented opportunities for improved situational awareness and understanding as well as ‘deep learning’ across these different silos.

The promise of AI for the more political (and thus emotional) D2G/G2G layer is less straightforward. AI is likely to provide our politicians as well as our citizens with uniquely deep insights in both the drivers of various types of human endeavors and their outcomes. The extent to which wily politicians will prove amenable to conform their behavior with these insights remains an open question.

4.2.4. Defense and Security Ecosystem: E2I (/E2E)

The fourth layer (which exists only latently in most developed societies) is much broader than the previous three, but it is in our view also potentially far more potent – also from the point of view of the applicability of AI. This layer includes both the level beyond government (the defense and security ecosystem) and how it interacts with the layer underneath all of these institutions (the individual). We have noted in previous work how the increased technology-driven empowerment of this most human of layers is one of the most powerful mega-trends in this day and age²²⁰ – and one that some of the West’s opponents appear to have been more adroit and adept at embracing/exploiting than we have.

Since the dawn of the industrial age our nation states have confided the legitimate use of (increasingly awesome) physical power to a uniformed corps of defense and security professionals, equipped with the requisite ‘arms’ to maintain their dominant position as the ultimate agents of power with the ability of physical escalation dominance. Just like the physical technologies that the industrial revolution spawned and that so radically transformed the manifestations of violence and war, so too did these industrial-age social technologies. Many recent authoritative long-term histories of political violence²²¹

220 Stephan De Spiegeleire et al., *Si Vis Pacem, Para Utique Pacem. Individual Empowerment, Societal Resilience and the Armed Forces* (The Hague, The Netherlands: The Hague Centre for Strategic Studies, 2015), <http://www.literatuurplein.nl/boekdetail.jsp?boekId=1078494>

221 Francis Fukuyama, *The Origins of Political Order: From Prehuman Times to the French Revolution* (Farrar, Straus and Giroux, 2011); Francis Fukuyama, *Political Order and Political Decay: From the Industrial Revolution to the Globalization of Democracy* (Farrar, Straus and Giroux, 2014); Azar Gat, *War in Human Civilization* (OUP Oxford, 2008); Yuval Noah

have documented how this monopolization of violence has greatly reduced the overall incidence of different types of violence, even though there clearly still remains ample room for improvement.


We see all around us, however, how new technologies are starting to transform our economies, societies and politics – in many cases quite disruptively. Our militaries will not remain unaffected. Like all purposive actors (public and private) they will have to start redesigning the way they interact with others in order to achieve the defense and security objectives of which they remain the legitimate custodians.

If we look at how this process is unfolding in the private sector, we observe the emergence of truly novel ways of organizing purposive action. Companies still exist; they still compete; they still want to prevail in their markets; they still want to make money. But they are starting to do all of this in radically different ways. The astonishing success of these new models may portend analogous changes in the military world. In *The Platform Revolution*, a popular recent book on this phenomenon, three influential authors describe how disruptors like Airbnb, Amazon, Alibaba, Facebook, Google, Uber, and many others have found new ways to leverage technology to connect people, organizations, and resources around themselves into interactive ecosystems in which “amazing amounts of value can be created and exchanged. Even incumbent giants from Walmart and Nike to John Deere, GE, and Disney are all scrambling to adopt similar models”.²²²

In some sense, this is a Copernican revolution in the way in which these private actors organize and position themselves in their environment to achieve their goals. The way they generate value is still through their own selfishly purposive efforts, but these efforts now increasingly put their ecosystem partners, and not themselves or their traditional supply chains, center stage. Rather than seeing themselves at the center of their business universe – doing innovation mostly on the inside; selecting and controlling supply chain partners; shielding themselves off from competitors – they are starting to see themselves as much less central (but therefore not less profitable) players who increasingly get to generate even more private value by piggy-backing on the efforts of outsiders – global innovative players that they do not choose, but that choose them.

Harari, *Homo Deus: A Brief History of Tomorrow* (Random House, 2016); Yuval Noah Harari, *Sapiens: A Brief History of Humankind* (Random House, 2014); Steven Pinker, *The Better Angels of Our Nature: Why Violence Has Declined* (Penguin, 2011); Ian Morris, *Why The West Rules - For Now: The Patterns of History and What They Reveal about the Future* (Profile Books, 2010); Johan Norberg, *Progress: Ten Reasons to Look Forward to the Future* (Oneworld Publications, 2017)

²²² Parker, Van Alstyne, and Choudary, *Platform Revolution*. On the emergence of new forms of organizing purposive action (what we called here ‘new social technologies’) outside of the military realm, see also Sangeet Paul Choudary, *Platform Scale: How an Emerging Business Model Helps Startups Build Large Empires with Minimum Investment*, First edition (S.L.: Platform Thinking Labs Pte. Ltd, 2015); David S. Evans and Richard Schmalensee, *Matchmakers: The New Economics of Platform Businesses: How One of the Oldest Business Models on Earth Powers the Most Incredible Companies in the World* (Boston, Massachusetts: Harvard Business Review Press, 2016); *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*, 2014; Salim Ismail and Michael S. Malone, *Exponential Organizations: Why New Organizations Are Ten Times Better, Faster, and Cheaper than Yours (and What to Do about It)*, 1. ed, A Singularity University Book (New York, NY: Diversion Books, 2014).




Unlike companies, defense organizations are not in the business of producing private financial value. They are, however, there to produce public security value. Recent developments in geodynamics (Russia, ISIS, terrorism, etc.) suggest that the demand for that defense security value is increasing – both by preventing and dealing with security threats, and by jumping on security opportunities²²³. The *en-vogue* term ‘hybrid warfare’ suggests that also some of our adversaries are experimenting with such new, less linear and more ‘piggybacking’ models for organizing their defense efforts.

The ‘platform ecosystem’ model might offer some lessons for a different way in which our defense organizations might provide defense and security value for their taxpayers’ money – re-actively against traditional and ‘new’ challenges, but also pro-actively to new reap new opportunities for prevention and building resilience-building. Many of the key trends that have enabled the platform ecosystem revolution in the private sector are equally relevant and visible in the security realm. Most of the emerging physical and, increasingly, digital technologies that connect and empower business ecosystems are organically intertwined with the security dynamics in our societies. Social networks connect both the agents of conflict and the agents of resilience; the agents of progress and the agents of regress. Just like companies can now increasingly leverage available market intelligence about individuals’ preferences ‘from the source’ – the (big) data that is being generated and collected – but also the individual creativity and drive of individual and networked entrepreneurs, so too could (legitimate and responsible) defense and security professionals find ways to leverage the mega-trend of global personal and inter-personal empowerment in order to achieve our defense and security objectives.

Our current myopic focus on the first three layers of ‘defense and security’ may have obscured the existence of an increasingly vibrant defense and security ecosystem ‘out there’. With many diverse actors – some known, many more unknown – that could help our defense and security organizations in achieving their goals, both as sensors and as effectors. This ecosystem may include app developers that are currently using big data to develop apps for ‘our’ daily needs, but that could also be incentivized to dream up apps that could address defense and security challenges in fragile, conflict-ridden or propaganda-afflicted regions. It may include individual youngsters, farmers, fishermen, mothers with smartphones living in those same regions, who could be empowered as agents of stability or resilience. It includes the rapidly exploding global deep learning and AI community that is coming up with new ways to diagnose and address various problems (in terms of economy, education, health, but also security) much earlier in the process. It certainly includes companies that benefit from stability in developing their markets, but that also contribute to the same through their investments.

Could our defense organizations start using big data and AI to become the catalysts and custodians of these broader defense and security ecosystems – not by trying to ‘command & control’ the ecosystem actors, but by stimulating, facilitating and leveraging them? Can

²²³ See the recent HCSS Strategic Monitors for more details De Spiegeleire and Sweijs, Volatility and Friction in the Age of Disintermediation. *HCSS StratMon 2016-2017 Yearly Report*; De Spiegeleire et al., *Si Vis Pacem, Para Utique Pacem. Individual Empowerment, Societal Resilience and the Armed Forces*.



we imagine defense capability planning processes that try to include not only their ‘own’ capabilities, but that start looking for a better balance between those ‘own’ capabilities and capabilities to empower various ecosystem players? Whereby ‘defense’ is no longer just the operator who intervenes heavily, physically, lately in a conflict dynamic (and we still foresee a need for that security function as well) ; but the more strategic custodian who looks for ways to intervene, digitally and early in the process? A thoughtful curator that advocates a better balance between conflict-centric and resilience-centric efforts in order to maximize the defense value proposition?

4.3. AI in Other Armed Forces Today

The armed forces of the world’s leading military powers all recognize the qualitative edge AI systems are likely to give them today and tomorrow – soldiers who often “face problems of scale, complexity, pace and resilience that outpace unaided human decision making.”²²⁴ Whether for commanders faced with unconventional adversaries in high-speed engagements; intelligence analysts faced with drawing the correct conclusions from petabytes of noisy data; or frontline soldiers; AI promises to augment analysis and decision-making capabilities and reaction times both, speed up learning, and improve their ability to act with discretion, accuracy, and care under uncertain and changing conditions. It is therefore little surprise that many of world’s leading militaries are running active AI development and deployment programs; however, it is the way these AI systems are tailored to underlying needs, which reveals a lot about the evolving strategic and tactical doctrines of these powers – and the changing nature of deterrence and warfare in the decades to come.

In this section, we present some evidence that can be gleaned from open sources on the military AI efforts made by four particularly active countries, which will be surveyed in purely alphabetic order.

4.3.1. China

As the second biggest ‘player’ in general-purpose AI China is increasingly showing that it is more than capable of keeping pace with the US in this field. While in terms of fundamental breakthroughs, China is still lagging behind the US, there has been a massive increase in growth in terms of cited (machine learning) research.²²⁵ To spur this, in February 2017, China’s National Development and Reform Commission approved a plan to establish an online ‘national laboratory for deep learning’, commissioning Baidu to set up the research effort which will focus on seven areas of research including machine learning-based visual recognition, voice recognition, new types of human

224 Artificial Intelligence · Lockheed Martin,” accessed September 21, 2016, <http://www.lockheedmartin.com/us/atl/research/artificial-intelligence.html>.

225 Zhang, “China’s Artificial-Intelligence Boom.”. Cf. Niu et al., “Global Research on Artificial Intelligence from 1990–2014.”

machine interaction and deep learning intellectual property. The overarching goal, it stated, is to “boost China’s overall competence in artificial intelligence”.²²⁶ Meanwhile, major Chinese companies such as Baidu, Alibaba and Tencent have proven remarkably adept at rapidly iterating over breakthroughs to develop – and deploy – applications of this technology,²²⁷ as well as making remarkable home-grown breakthroughs in fields such as speech recognition²²⁸ or self-driving cars.²²⁹

These developments are catalyzed by a high-performance computing (HPC) industry which is increasingly self-reliant: after the US government banned the sale of powerful Intel Xeon processors to Chinese supercomputing initiatives in April 2015,²³⁰ China was able to substitute its own, native-built processors in the design of the Sunway TaihuLight, since 2016 the world’s fastest supercomputer.²³¹

Moreover, this drive towards AI is spurred on by strong links between private actors and civilian applications on the one hand, and government agencies (specifically the People’s Liberation Army - PLA) on the other. Such government support goes all the way to the top: in 2016, China’s 13th 5-year plan (2016-20) highlighted the importance of further breakthroughs in AI, as did the 13th 5-year National Science and Technology Innovation Plan.²³² In 2016, the Chinese government announced plans to develop a 100 billion RMB (\$15 billion) AI market by 2018.²³³ These initiatives have been characterized as part of the “China Brain Plan” (中国脑计划),²³⁴ an ambitious effort to develop artificial intelligence and deploy it in unmanned systems, in cyber security and for social governance – and for military supremacy – under the umbrella of President Xi Jinping’s national strategy of “military-civil fusion” (军民融合).²³⁵

At present, the Chinese military’s initial approach to artificial intelligence is still strongly informed by its examination of US developments and initiatives. This echoes historical context: in the wake of the 1st Gulf War, as US forces demonstrated the supremacy of

226 Meng Jing, “China’s First ‘deep Learning Lab’ Intensifies Challenge to US in Artificial Intelligence Race,” *South China Morning Post*, February 21, 2017, <http://www.scmp.com/tech/china-tech/article/2072692/chinas-first-deep-learning-lab-intensifies-challenge-us-artificial>.

227 Zhang, “China’s Artificial-Intelligence Boom.”

228 Will Knight, “Baidu System Rivals People at Speech Recognition,” *MIT Technology Review*, 2015, <https://www.technologyreview.com/s/544651/baidus-deep-learning-system-rivals-people-at-speech-recognition/>.

229 Aaron Mamiit, “China’s Baidu Unveils New All-Electric, Self-Driving Car: Testing Begins For Modified Chery EQ,” *Tech Times*, August 28, 2016, <http://www.techtimes.com/articles/175317/20160828/chinas-baidu-unveils-new-all-electric-self-driving-car-testing-begins-for-modified-chery-eq.htm>

230 Don Clark, “U.S. Agencies Block Technology Exports for Supercomputer in China,” *Wall Street Journal*, April 9, 2015, sec. Tech, <http://www.wsj.com/articles/u-s-agencies-block-technology-exports-for-supercomputer-in-china-1428561987>

231 Zhang, “China’s Artificial-Intelligence Boom.”

232 People’s Republic of China, “The People’s Republic of China National Economic and Social Development of the Thirteenth Five-Year Plan,” 2016, http://www.gov.cn/xinwen/2016-03/17/content_5054992.htm.; People’s Republic of China, “13th Five-Year National Science and Technology Innovation Plan,” 2016, http://www.gov.cn/zhengce/content/2016-08/08/content_5098072.htm

233 Xiao, “The Chinese Government Wants A 100 Billion RMB AI Market By 2018.”

234 Xinhua News Agency, “中国‘脑计划’纳入规划全面展开 坚持‘一体两翼,’” Xinhua News Agency, 2016, http://news.xinhuanet.com/politics/2016-08/18/c_129238381.htm.; Elsa Kania, “China May Soon Surpass America on the Artificial Intelligence Battlefield,” Text, *The National Interest*, (February 21, 2017), <http://nationalinterest.org/feature/china-may-soon-surpass-america-the-artificial-intelligence-19524>

235 Kania, “China May Soon Surpass America on the Artificial Intelligence Battlefield.”

network-centric warfare, the PLA embarked on an ambitious agenda of ‘informatization’ (信息化).²³⁶ While it has made great strides in this space, the PLA ultimately has not yet caught up to evolving US integration of information technologies. As such, being faced with the US’s ‘Third Offset Strategy’, the PLA has identified a great need – and opportunity – to use the disruptive potential of AI to match or even leapfrog US capabilities, seeking what it calls a “military revolution of intelligentization”: moving beyond “digitalization” (数字化) and “networkization” (网络化), to achieve the transition from ‘informatized’ to ‘intelligentized’ (智能化) warfare.²³⁷

Concretely, as part of this effort the PLA has established an Intelligent Unmanned Systems and Systems of Systems Science and Technology Domain Expert Group (军委智能无人系统及体系科学技术领域专家组), and has invested in and field-tested intelligent unmanned vessels for reconnaissance missions and to reinforce the PLA Navy’s ability to monitor and establish a presence in disputed waters.²³⁸ The PLA has also made breakthroughs in UAV swarming and command and control, and even – inspired by AlphaGo’s victory over Lee Sedol in the ancient Chinese game of go – explored the role of intelligentized command and control within a joint operations command system.²³⁹

Another area of Chinese interest is missile technology: in order to counter the US’s ‘semi-autonomous’ Long Range Anti-Ship Missile (LRASM) – a replacement for the Harpoon missile, capable of autonomously following waypoints, avoiding engagement range of non-target ships, and optimizing strike location for maximum lethality²⁴⁰ – China has expressed interest in utilizing AI to empower the flight guidance and target recognition systems in new generations of its own cruise missiles. This could enhance their operational versatility, allowing commanders to tailor missiles to specific and rapidly altering battlefield conditions.²⁴¹ Wang Changqing, director of the General Design Department of the China Aerospace Science and Industry Corp’s 3rd Academy, told China Daily that Chinese engineers have researched such applications for AI for many years, and they are leading the world, adding that “our future cruise missiles will have a very high level of artificial intelligence and automation. They will allow commanders to control them in a real-time manner, or to use a fire-and-forget mode, or even to add more tasks to in-flight missiles.”²⁴² This, Chinese officials hope, may enable and strengthen a doctrine

236 Elsa Kania, “The Next U.S.-China Arms Race: Artificial Intelligence?,” *The National Interest*, 2017, <http://nationalinterest.org/feature/the-next-us-china-arms-race-artificial-intelligence-19729>

237 Ministry of National Defense of the People’s Republic of China, “The Dawn of the Intelligent Military Revolution” (People’s Liberation Army Daily, January 28, 2016), http://www.mod.gov.cn/wqzb/2016-01/28/content_4637961.htm.; Kania, “The Next U.S.-China Arms Race.”

238 Kania, “China May Soon Surpass America on the Artificial Intelligence Battlefield.”

239 Ibid.; Kania, “The Next U.S.-China Arms Race.”

240 Lockheed Martin, “LRASM: Overview,” *Lockheed Martin*, 2016, <http://www.lockheedmartin.com/us/products/LRASM/overview.html>

241 Zhao Lei, “Nation’s next Generation of Missiles to Be Highly Flexible,” *China Daily*, 2016, http://www.chinadaily.com.cn/china/2016-08/19/content_26530461.htm.; Ben Blanchard, “China Eyes Artificial Intelligence for New Cruise Missiles,” *Reuters*, August 19, 2016, <http://www.reuters.com/article/us-china-defence-missiles-idUSKCN10U0EM>.; Abhijit Singh, “Is China Really Building Missiles With Artificial Intelligence?,” *The Diplomat*, 2016, <http://thediplomat.com/2016/09/is-china-really-building-missiles-with-artificial-intelligence/>

242 Lei, “Nation’s next Generation of Missiles to Be Highly Flexible.”; Blanchard, “China Eyes Artificial Intelligence for New Cruise Missiles.”

known as 'remote warfare', whereby large fleets of small vessels are able to successfully attack, and evade or overwhelm the point defense systems of larger capital ships such as US aircraft carriers.²⁴³

Significantly, given that the PLA conventionally approaches military innovation through a lens of 'technology determines tactics',²⁴⁴ they may be more inclined to rely upon artificial intelligence than the US – and may be more willing to relinquish 'meaningful human control' in order to achieve ever-greater cognitive speed in battlefield decisions.

4.3.2. Israel

Israel was one of the first countries to reveal that it has deployed fully automated robots: self-driving military vehicles to patrol the border with the Palestinian-governed Gaza Strip. Next in the IDF's plans is to equip the vehicles with weapons, and deploy them in stages to Israel's frontiers with Egypt, Jordan, Syria, and Lebanon. Meanwhile, the Israeli 'Harpy' anti-radiation unmanned aerial vehicle is claimed to already be able to detect, target, and engage enemy radar installations without any human oversight or supervision.²⁴⁵ Further in the future, the military is looking to form mixed combat units of robotic vehicles and human soldiers.²⁴⁶

Various Israeli companies also claim to apply AI in a number of their defense systems. The former chief of the Israel Defense Forces' (IDF's) Central Command and current head of land systems at Israel Aerospace Industries (IAI), Major General (ret'd) Gadi Shamni, for instance, told IHS Jane's that the company has spent enormous sums on research and development (R&D) aimed at developing new products to enable the IDF and overseas clients to detect and accurately strike time-critical targets. "In an age of big data, this requires a system that can handle such heavy information loads so IAI has developed an automated system, called Automated Decision Making (ADM), that employs artificial intelligence and robotics to sift through the data and respond to it instantly."²⁴⁷

Israeli contractor Aeronautics Ltd has also produced a range of UAV control systems which allegedly contain AI-algorithms. One of those systems, Aeronautics' Unmanned

243 Cf. John Markoff and Matthew Rosenberg, "China's Intelligent Weaponry Gets Smarter," *The New York*

Times, February 3, 2017, <https://www.nytimes.com/2017/02/03/technology/artificial-intelligence-china-united-states.html>

244 Cf. Dennis J. Blasko, "Technology Determines Tactics': The Relationship between Technology and Doctrine in Chinese Military Thinking," *Journal of Strategic Studies* 34, no. 3 (June 1, 2011): 355–81, doi:10.1080/01402390.2011.574979

245 Paul Scharre, "Robotics on the Battlefield, Part I: Range, Persistence and Daring," 20YY Series: Preparing for War in the Robotic Age. (Center for a New American Security, May 2014), http://www.cnas.org/sites/default/files/publications-pdf/CNAS_RoboticsOnTheBattlefield_Scharre.pdf, footnote 50

246 The Future of War: Israel First to Deploy Fully Automated Military Robots," *The Mainichi*, August 24, 2016, <http://global.factiva.com/redir/default.aspx?P=sa&an=AIWMDM0020160824ec8000105&cat=a&ep=ASE>; Israel Today, "Israel First to Deploy Autonomous Military Robots," *Israel Today*, August 25, 2016, <http://www.israeltoday.co.il/NewsItem/tabid/178/nid/29924/Default.aspx>.

247 Yaakov Lappin, "Interview: Gadi Shamni, Head of Land Systems, Israel Aerospace Industries," *Jane's Defence Weekly*, September 7, 2016, <http://janes.ihs.com.proxy1.athensams.net/Janes/Display/1782846>

Multi-Application System™ (UMAS), is “a software-based package that is designed to provide ‘advanced’ control of a ‘variety’ of manned and unmanned applications. As such, it is described as incorporating proprietary artificial intelligence and ‘unique’ interfaces and as offering ‘unrivalled’ levels of system reliability and performance. System functions include artificial intelligence-based electronic/mechanical failure prediction; real-time decision making support; the facilitating of data transfer between systems; and hierarchical handling of fixed/variable order data strings as a function of tool safety (using an integral time domain camouflage interface).”²⁴⁸

Also in the field of combat simulation, Israeli defense electronics company Elbit Systems Ltd produces a Command and Staff Trainer (CST) that simulates a range of joint operations. The system is purported to include advanced simulation models (full spectrum of conflict, mounted operations in urban terrain, out-of-the-window homeland security) that contain a scenario with thousands of entities and large training areas and include artificial intelligence for aggregate behaviors.²⁴⁹ The system was delivered to the Royal Netherlands Army (RNLA) in 2012 for use in command and staff exercises, for concept development and experimentation (CD&E), as well as mission rehearsal events to simulate complex battlefield operations in both low and high-intensity conflicts as well a broad spectrum of non-military settings, such as operations other than war (OOTW) and civil-military cooperation (CIMIC), including operation of military forces with municipal authorities, police and medical organization.²⁵⁰

As is increasingly the case in the civilian world, AI is often used as a sales pitch. It is virtually impossible to independently verify the various claims made by these companies. Israel’s prowess in AI is widely acknowledged, however, also stimulated by the many startups that emerge in the IDF’s ecosystem.²⁵¹

4.3.3. Russia

While still somewhat lagging behind on its great power rivals in terms of deep machine learning capabilities, the Russian Federation has displayed a steady commitment to developing and deploying a wide range of robotic military platforms, including unmanned ground vehicles (UGVs), with the full backing of its MoD and domestic industries: in January 2017, President Putin called for the creation of “autonomous robotic complexes”

248 Jane’s, “Aeronautics Control Systems,” *Jane’s Unmanned Aerial Vehicles and Targets*, 22-Mar-17, <http://janes.ihc.com.proxy1.athensams.net/Janes/Display/1318450>

249 Jane’s, “Elbit Systems Ltd - Command and Staff Training Systems,” *Jane’s Simulation and Training Systems*, March 23, 2017, <http://janes.ihc.com.proxy1.athensams.net/Janes/Display/1592554>.

250 army-technology.com, “Elbit Delivers Command and Staff Trainer to Royal Netherlands Army - Army Technology,” November 23, 2012, <http://www.army-technology.com/news/newselbit-delivers-command-staff-trainer-royal-netherlands-army>

251 Ariel Felner, “The Israeli AI Community,” *AI Magazine* 37, no. 3 (October 7, 2016): 118–22; Eze Vidra, “30 Machine Intelligence Startups to Watch in Israel,” *Medium*, February 15, 2017, <https://medium.com/@ediggs/30-machine-intelligence-startups-to-watch-in-israel-a05b6597c4a5>. For more on the Israeli start-up culture See also Dan Senor and Saul Singer, *Start-up Nation: The Story of Israel’s Economic Miracle*, Reprint edition (New York: Twelve, 2011).

for use by the military,²⁵² previously entrusting the Advanced Research Foundation with the new 'National Center for the Development of Robotic Technologies and Basic Robotic Components', in order to consolidate so far uncoordinated efforts for the creation of advanced robotic equipment. Andrey Grigoryev, General Director of the Foundation, has noted that "the centre will allow for efficiently sharing information about already available developments and achievements, will lead to their unification, to making the world's best military-use and special-purpose robots from the best components, including control systems, visual equipment complexes, special sensors, and much more. We have many rather interesting domestic achievements in the most advanced sectors."²⁵³ In an interview by the official Russian newspaper *Rossiyskaya Gazeta* in August 2016, Grigoryev detailed an exercise in the Nizhniy Tagil training range in which an unmanned helicopter quickly detected a concealed 'enemy', issued a command to a tracked 'terminator-robot' armored vehicle named 'Nerekhta' armed with missiles and machine guns that deployed to a firing position, made its own decisions about what would be best to use in a given situation to destroy the targets, and accomplished the mission of destroying the 'enemy' successfully.²⁵⁴

More critically, while most Russian robotic UGVs are still in development and testing stages, a number of them have actually seen active combat service: the Uran-6 demining robot saw use by Russian forces operating in Syria, amongst others for clearing Palmyra of booby traps and IEDs left behind by ISIS forces in the wake of the March 2016 regime offensive against the city.²⁵⁵ Already, Russian defense experts have suggested that the robot's larger cousin, the heavily armed and armored Uran-9, may be deployed in Syria in support of Russian or Syrian regime ground operations,²⁵⁶ although there are indications that to present this has not yet taken place.²⁵⁷ Similarly, the 'Platforma-M' reconnaissance UGV, developed by the 'Progress' Science and Technical Institute, is currently deployed with the Russian Pacific Fleet.²⁵⁸ The heavy 'Udar' UGV, unveiled in 2015, comes in combat, engineering support, and transportation versions, and has been built on the existing frame of the BMP-3 armored vehicle, in order to ease the maintenance and repair of the system.²⁵⁹

252 ФБА «Экономика сегодня», "Назад В Будущее: Путин Объявил Эру Боевых Роботов," *Рамблер/Новости*, January 26, 2017, <https://news.rambler.ru/politics/35932393-nazad-v-buduschee-putin-obyavil-eru-boevyih-robotov/>.

253 Сергей Птичкин, "Умная пуля видит цель," *Российская Газета*, July 19, 2016, sec. СОБЫТИЯ И КОММЕНТАРИИ; Арсенал

254 Russian Pundit Interviewed on Latest Combat Equipment," *BBC Monitoring Former Soviet Union*, August 4, 2016, <http://global.factiva.com/redir/default.aspx?P=sa&an=BBCSUP0020160804ec8400m9&cat=a&ep=ASE>

255 Russia Today, "Russia's Mine-Clearing Uran-6 Robots to Help Get Rid of Hidden Explosives in Palmyra (VIDEOS)," *RT International*, March 30, 2016, <https://www.rt.com/news/337810-russia-palmyra-demine-robot/>; cf. Samuel Bendett, "Get Ready, NATO: Russia's New Killer Robots Are Nearly Ready for War," Text, *The National Interest*, (March 7, 2017), <http://nationalinterest.org/blog/the-buzz/russias-new-killer-robots-are-nearly-ready-war-19698>.

256 Эксперт Не Исключил, Что Новейший комплекс 'Уран-9' испытают В Сирии," *РИА Новости*, January 9, 2017, <https://ria.ru/arms/20170109/1485307092.html>; Ян Грогман, "Российский Робот-Танк «Уран-9»: Бесплезная Игрушка Или Революция?," *ИноСМИ.Ru*, April 13, 2016, <http://inosmi.ru/military/20160413/236108861.html>

257 Арик Толер, "Использовала Ли Россия Боевых Роботов В Сирии?," *ИноСМИ.Ru*, January 17, 2016, <http://inosmi.ru/military/20160117/235078297.html>

258 *nortwolf_sam*, "Роботизированный Комплекс «Платформа-М» На Вооружении Тихоокеанского Флота," *Nortwolf_sam*, October 11, 2015, <http://nortwolf-sam.livejournal.com/1054499.html>

259 Bendett, "Get Ready, NATO."

Notably, at present most of these Russian military bots are still designed (or deployed) to be controlled by remote operator, as the Russian defense establishment remains somewhat uncomfortable with the notion of fully autonomous military systems:²⁶⁰ in 2015, the Russian daily *Komsomolskaya Pravda* still argued that “such [fully autonomous] weapons are of an offensive nature, while our military doctrine is defensive [...] If there is an artificial brain in such a system, you never know what may happen, since ‘friend or foe’ computer and signal mechanics can be easily suppressed by means of electronic warfare, which is one of Russia’s key military strengths.”²⁶¹

However, while Russia might put some faith in its electronic warfare abilities to try and scramble or nullify the autonomous weapons deployed by its adversaries, it is possible this attitude will change should the Russian military face increasing numbers of fully autonomous, and highly responsive combat systems.²⁶²

Finally, Russia’s United Instrument Manufacturing Corporation (OPK) is working on using AI systems for border protection, developing a system which will automatically interact with video cameras, infrared and seismic sensors, radars and drones, in order to monitor and observe any type of violations. In addition, the new system, to be deployed on Russia’s Eastern and Southern borders, is intended not only to collect different types of information, but also contains elements of artificial intelligence which will allow for analysis and forecasting of the situation and work out proposals for the protection of borders, by calculating steps and routes that offenders may take, as well as the necessary measures to prevent malicious acts, including the assessment of possible risks.²⁶³

4.3.4. US

As of yet the most prominent actor in the field of military AI, the United States has been actively involved in AI-related R&D since its very emergence of the field (see chapter 2 page 26). However, having overseen and driven much of the early breakthroughs in AI research (and computer science broadly) during the Cold War, the bulk of the US’s efforts have now shifted towards deploying these technologies with a clear focus on increasing its armed forces’ operational effectiveness as well as standoff force projection capabilities (A2A in our terminology).

One well-known driver of the US’s commitment to developing these systems is their extensive military experience, over the last two decades, with drone systems. As a result of this experience, notably in Iraq and Afghanistan, US spending on Unmanned Aircraft Systems (UAS) grew tenfold, from \$283 million in 2000 to \$2.9bn in 2016. At the same time, the US inventory of UAS exploded a staggering 65x (from 167 to 11,000) between

260 Samuel Bendett, “Can Russia’s Military Bots Keep Pace?,” *RealClearWorld*, August 27, 2015, http://www.realclearworld.com/blog/2015/08/robots_military_russia_hopes.html; Bendett, “Get Ready, NATO.”

261 Михаил ТИМОШЕНКО | Сайт «Комсомольской правды», “Война Под Микроскопом,” *KP.RU - Caйм «Комсомольской Правды»*, August 6, 2015, <http://www.kp.ru/daily/26415/3289072/>.

262 Bendett, “Get Ready, NATO.”

263 Sputnik, “Robots on Patrol: Russian Borders to Be Guarded by Artificial Intelligence,” 2016, <https://sputniknews.com/military/201606301042223143-russian-borders-guarded-robots/>.

2002 and 2013, even in the face of multi-year downturns in overall US defense spending.²⁶⁴ The patent success of these systems has as such led to the formulation of the ‘Unmanned Systems Integrated Roadmap’, setting out an ambitious roadmap for further developing all Unmanned Autonomous Systems (airborne, ground and maritime) in the period up to 2038.²⁶⁵

However, US ambitions in the field of ‘intelligent weaponry’ go far beyond this; in 2015, the Pentagon’s fiscal 2017 budget request included \$12-15 billion to fund war gaming and the demonstration of new technologies – including wearable electronics, exoskeletons, autonomous weapons and unmanned aircraft; drone mother ships and deep-learning machines – which could ensure a continued military edge over great powers such as China and Russia.²⁶⁶ These investments, and the major role they reserve for AI in future military force projection, reflect the core logic of the Pentagon’s so-called ‘Third Offset Strategy’.²⁶⁷

This doctrine, in development since 2012, was first announced in November 2014 by Chuck Hagel,²⁶⁸ and later pursued and implemented by Ashton B. Carter, both successive US Secretaries of Defense under President Obama.²⁶⁹ Aiming at a geostrategic doctrine that can preserve a qualitative military advantage in the face of rising rivals,²⁷⁰ the ‘Third Offset’ strategy is intended to be the successor to the First Offset Strategy of the 1950s – in which President Dwight D. Eisenhower sought to develop nuclear-warhead and missile technologies to counter and deter Soviet conventional numerical superiority – and the Second Offset Strategy of the 1970s – in which Secretary of Defense Harold Brown shepherded the development of precision-guided munitions, stealth, and intelligence, surveillance, and reconnaissance (ISR) systems to counter the nuclear parity and numerical superiority and improving technical capability of Warsaw Pact forces along the Central Front in Europe.²⁷¹

264 Bank of America Merrill Lynch, “Robot Revolution - Global Robot & AI Primer.”, pg. 7

265 Department of Defense, “Unmanned Systems Integrated Roadmap: FY2013-2038,” 2013, <http://archive.defense.gov/pubs/DOD-USRM-2013.pdf>.

266 Shalal, “Pentagon Eyes \$12-15 Billion for Early Work on New Technologies.”; cf. also Mackenzie, “The Future Military-Artificial Intelligence Complex? | FT Alphaville,” *Financial Times* - Alphaville, 2015, <http://ftalphaville.ft.com/2015/12/15/2147846/the-future-military-artificial-intelligence-complex/>.

267 Joshua Pavluk and August Cole, “From Strategy to Execution: Accelerating the Third Offset,” *War on the Rocks*, June 9, 2016, <http://warontherocks.com/2016/06/from-strategy-to-execution-accelerating-the-third-offset/>. John Markoff, “Pentagon Turns to Silicon Valley for Edge in Artificial Intelligence,” *The New York Times*, May 11, 2016, <http://www.nytimes.com/2016/05/12/technology/artificial-intelligence-as-the-pentagons-latest-weapon.html>.

268 Cf. Robert Martinage, “Toward a New Offset Strategy: Exploiting U.s. Long-Term Advantages to Restore U.s. Global Power Projection Capability” (*Center for Strategic and Budgetary Assessments*, 2014), <http://csbaonline.org/uploads/documents/Offset-Strategy-Web.pdf>.

269 Ashton Carter, “Keynote Address: The Path to the Innovative Future of Defense,” *Center for Strategic and International Studies*, October 28, 2016, https://csis-prod.s3.amazonaws.com/s3fs-public/event/161028_Secretary_Ashton_Carter_Keynote_Address_The_Path_to_the_Innovative_Future_of_Defense.pdf. The authors of this report recognize that the change in administration which has occurred since the articulation of the Third Offset has re-opened many settled strategic assumptions, and may lead to major changes in the US force posture, operational doctrine, and procurement policies. Nonetheless while some modest structural changes may occur, there is at present no reason to suppose a drastic change in the role of AI within the doctrine. Cf. also Theodore R. Johnson, “Donald Trump’s Pentagon and the Future of the Third Offset Strategy: Will the Department of Defense Invest in People or Technology?,” *The Atlantic*, November 29, 2016, <https://www.theatlantic.com/politics/archive/2016/11/trump-military-third-offset-strategy/508964/>.

270 Markoff and Rosenberg, “China’s Intelligent Weaponry Gets Smarter.”

271 Timothy A. Walton, “Securing The Third Offset Strategy: Priorities For Next US Secretary Of Defense – Analysis

What does that look like in practice? Speaking in 2015, Robert Work, the then-US deputy secretary of defense, emphasized “human-machine collaboration combat teaming”, arguing that: “Early adoption will be a key competitive advantage, while those that lag in investment will see their competitiveness slip”.²⁷² In this speech to the Defense One National Security Forum conference, Work identified five pillars to the military future:²⁷³

1. **Autonomous deep learning machine systems** which are able to see the patterns through the chaff of hybrid warfare, to give early warning that something is happening in gray zone conflict areas (such as the Ukraine), and which are able to respond at extreme speed, and under rapidly shrinking engagement windows. Such learning systems might, he argues, fill the gap in those fields – such as air defense or cyber defense – where human operators alone cannot achieve sufficient speed to stop or degrade a determined attack.
2. **Human machine collaboration**, which will include the promotion of so-called ‘Centaur’ warfighting,²⁷⁴ going from the observation that teams combining the strategic analysis of a human with the tactical acuity of a computer, reliably defeat either human-only or computer-only teams across many games.
3. **Assisted human operations**, whereat wearable electronics, uploadable combat apps; heads up displays, exoskeletons, and other systems, can enable humans on the front line to perform better in combat.
4. **Advanced human-machine combat teaming** where a human working with unmanned systems is able to take better decisions and undertake cooperative operations. Examples of these are the Army’s Apache and Gray Eagle UAV systems, which are designed to operate in conjunction. Other examples are drone ‘motherships’; electronic warfare networks, or swarming systems which will help transform operations by enabling one mission commander to direct a full swarm of micro-UAVs.
5. **Network-enabled semi-autonomous weapons**, where systems are both linked, and hardened to survive cyberattack.²⁷⁵

All of these systems would be networked together through learning systems, enabling a form of ‘algorithmic warfare’ and a machine-learning approach to targeting.²⁷⁶ In this

| CSBA, “CSBA | Center for Strategic and Budgetary Assessments, July 27, 2016, <http://csbaonline.org/2016/07/27/securing-the-third-offset-strategy-priorities-for-next-us-secretary-of-defense-analysis/>.. For a view on what this means for Europe, see Daniel Fiott, “Europe and the Pentagon’s Third Offset Strategy,” *The RUSI Journal* 161, no. 1 (January 2, 2016): 26–31, doi:10.1080/03071847.2016.1152118.

272 Mackenzie, “The Future Military-Artificial Intelligence Complex? | FT Alphaville.”

273 Ibid

274 Sydney J. Freedberg, “Centaur Army: Bob Work, Robotics, & The Third Offset Strategy,” *Breaking Defense*, November 9, 2015, <http://breakingdefense.com/2015/11/centaur-army-bob-work-robotics-the-third-offset-strategy/>

275 Mackenzie, “The Future Military-Artificial Intelligence Complex? | FT Alphaville.”

276 Cf. Charlie Lewis, “Capturing Flying Insects: A Machine Learning Approach to Targeting,” *War on the Rocks*, September 6, 2016, <http://warontherocks.com/2016/09/capturing-flying-insects-a-machine-learning-approach-to-targeting/>.. Note that such ‘algorithmic warfare’ also may call for oversight and accountability within the chain of command Cf. Dustin A. Lewis, Gabriella Blum, and Naz K. Modirzadeh, “War-Algorithm Accountability,” SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, August 31, 2016), <http://papers.ssrn.com/abstract=2832734>.

way, military operations can increase in accuracy and pace, and learning ability, all the while being guided by intelligence operations which are, through deep learning-enabled 'anticipatory intelligence', increasingly able to anticipate the development of social unrest and societal instability several days in advance.²⁷⁷

What is the doctrine backing up this strategic interest and investment? In terms of strategic emphasis, one of the best articulations of current US thinking can probably be found back in the 2015 study on Autonomy that was recently published (after a one-year delay) in June 2016 by the DOD's Defense Science Board (DSB)²⁷⁸, a committee of civilian experts that provides the Secretary of Defense with independent advice on scientific and technical matters²⁷⁹. The study focuses on "institutional and enterprise strategies to widen the use of autonomy; approaches to strengthening the operational pull for autonomous systems; and an approach accelerate the advancement of the technology for autonomy applications and capabilities".²⁸⁰ The officially commissioned but independent study concluded that "action is needed in all three areas to build trust and enable the most effective use of autonomy for the defense of the nation".

The focus of the study in terms of the areas where the DSB feels that autonomy can deliver value by mitigating operational challenges lies on: rapid decision-making; high heterogeneity and/or volume of data; intermittent communications; high complexity of coordinated action; danger of mission and High persistence and endurance. All of these challenges are quite (sensibly) generic, and might just as well be applied to preventing conflict (as opposed to fighting it) or to strengthening resilience. Yet the way in which they are operationalized in the study puts them all within the current focus on (mostly industrial-kinetic, but increasingly also cyber) operations that are predominantly focused on combatting the enemy. The heart of the study lies in its chapter 4 [page 61](#), which focuses on 'operational pull'.²⁸¹ Most of the concrete actions they endorse lie in the traditional realm – with a few cyber exceptions interspersed. In essence, this is an agenda to 'insert' AI into the currently dominant mode of thinking about armed force.

The study concludes with the observations that advances in artificial intelligence have ensured autonomy has now crossed a 'tipping point', and that such autonomous capabilities are increasingly and readily available to allies and adversaries, on the basis of

277 Frank Konkel, "The CIA Says It Can Predict Social Unrest as Early as 3 to 5 Days Out," *Defense One*, 2016, <http://www.defenseone.com/technology/2016/10/cia-says-it-can-predict-social-unrest-early-3-5-days-out/132121/>

278 Defense Science Board, "Autonomy."

279 Defense Science Board, "Defense Science Board - Charter," May 6, 2016, <http://www.acq.osd.mil/dsb/charter.htm>

280 Defense Science Board, "Defense Science Board Summer Study on Autonomy," 2016, <https://www.hsdl.org/?abstract&did=794641>.

281 The key passage here, from our point of view, is the following: "Because the DoD mission is so broad, it was beyond the scope of this study to conduct an exhaustive review of, and search for, all of the beneficial roles for autonomy. Rather, the study chose to select representative system and mission applications to illustrate the potential value of autonomy. The study investigated four areas in depth—protection, battlespace awareness, force application, and logistics. These are joint capability areas that could immediately adopt existing autonomous technologies." This illustrates nicely what we see as the profoundly disruptive nature of AI on the one hand and the current mode of operating in defense. The DSB study, while clearly sympathizing with the former, ends up on the side of the latter. This study advocates starting from the former.

which it recommends that the DoD “take immediate action to accelerate its exploitation of autonomy while also preparing to counter autonomy employed by adversaries.”²⁸²

The authors of this study have no quarrel with most of the recommendations of this DSB study. We do, however, feel that the study misses an important part of what AI can mean for defense and security from a strategic – and not merely operational – point of view. Our own hunch is that AI (and a number of attendant technological developments that are co-emerging around big data) may have a much more disruptive impact on the essence ‘defense’ than the focus on AI-enhanced physical robotics and how they might affect our current way of safeguarding defense suggest. This is what we want to focus on in the subsequent chapters.

4.4. Use Cases

This section provides a number of concrete use cases for AI in defense organizations. This list illustrates applications of AI useful to military purposes are not limited to solely enhancing ‘kinetic’ or ‘hard power’ functions in a tactical context, but also include supportive, logistical, and strategic applications providing forces with a qualitative edge, information, and staying power. The list reflects the current literature’s – as well as the leading nations’ defense (see section 4.3) – focus on what we have called the A2A layer, but also includes a number of use cases from the other layers where HCSS suspects better investment opportunities. As such, this overview is intended to be illustrative rather than exhaustive, and to spark debate and exploration, rather than settle it.

4.4.1. Automating Cyber Operations (G2G, D2D)

In the first place, AI systems can play a powerful role in their ‘native’ environment – cyberspace. There are indications that AI’s ability to sift through prohibitively large quantities of data, and pick up on vague cues can strengthen security: for instance, Distil Networks uses machine learning algorithms to begin to defend against Advanced Persistent Threat bots, whose interactions are normally difficult to discern from real human users;²⁸³ likewise, Google²⁸⁴ and the behavior analytics company Gurukul²⁸⁵ each separately have developed artificial intelligence-based approaches to cybersecurity based on user and risk profiles instead of human-defined rules. Scientists at Arizona State University have developed machine learning algorithms for identifying zero-day

282 Defense Science Board, “Autonomy.”, pg. lii; (cf. pg. 99-101)

283 Connor Forrest, “How Distil Networks Uses Machine Learning to Hunt down ‘Bad Bots’ - TechRepublic,” *TechRepublic*, August 3, 2016, <http://www.techrepublic.com/article/how-distil-networks-uses-machine-learning-to-hunts-down-bad-bots/>.

284 Cage Metz, “Google’s Training Its AI to Be Android’s Security Guard | WIRED,” *WIRED*, May 2, 2016, <http://www.wired.com/2016/06/googles-android-security-team-turns-machine-learning/>.

285 Mark Cox, “Machine-Learning Based Risk Analytics Vendor Gurukul Launches First Channel Program,” *ChannelBuzz.ca*, June 1, 2016, <http://www.channelbuzz.ca/2016/06/machine-learning-based-risk-analytics-vendor-gurukul-launches-first-channel-program-17747/>.

security exploits, and tracing them as they spread around the hacker community;²⁸⁶ and researchers at MIT's *Computer Science and Artificial Intelligence Laboratory* have worked with startup PatternEx to construct a ML system that can review more than 3.6 billion lines of log files each day, to detect 85% of attacks, autonomously learning from- and countering cyberattacks as they evolve in real time.²⁸⁷ Nonetheless, such systems can also be a double-edged sword, as they can facilitate the automation of vulnerability-detection and exploitation software. This was shown by the victory of MAYHEM, an AI developed by Carnegie Mellon University, in the 2016 Cyber Grand Challenge, a competition started by DARPA to spur the development of automatic cyber defense systems which can discover, prove and (if used defensively) correct software flaws and vulnerabilities in real time.²⁸⁸

4.4.2. Algorithmic Targeting (A2A)

Another major use case for AI, of particular value in a tactical context, is the use of AI in developing rapid and accurate automatic target recognition (ATR) systems: start-up Deep Learning Analytics, for instance, has developed a machine-learning based ATR program prototype for DARPA, to trial systems assisting pilots in finding and engaging targets.²⁸⁹ This is part of DARPA's Target Recognition and Adaptation in Contested Environments (TRACE) research program, which seeks to deliver an accurate, real-time, and energy-efficient target recognition system that can work with existing radar systems, to provide long-range targeting capabilities for tactical airborne surveillance.

This is useful since visual identification of targets by human pilots requires flying at close approach, putting aircraft at risk from anti-aircraft systems; conversely, although radar in principle enables the identification and engaging of ground targets at a standoff distance, this comes at the cost of unacceptably high false-alarm rates or collateral damage. ATR systems can thus combine the best of both worlds: highly accurate aircraft fire support from a standoff distance.²⁹⁰ Key tactical requirements for such a capability are that it can offer low false-alarm rates in complex environments – ensuring it is not easily thrown off by decoys, or mistakes – and can rapidly improve its learning capabilities on the basis of sparse or limited training data.²⁹¹

286 Eric Nunes et al., "Darknet and Deepnet Mining for Proactive Cybersecurity Threat Intelligence," *arXiv Preprint arXiv:1607.08583*, 2016, <http://arxiv.org/abs/1607.08583>

287 Kalyan Veeramachaneni and Ignacio Araldo, "AI2: Training a Big Data Machine to Defend," accessed August 20, 2016, http://people.csail.mit.edu/kalyan/AI2_Paper.pdf. For a shallow overview of startups and actors developing AI cybersecurity solutions, see AI.Business, "Artificial Intelligence in Defence and Security Industry," *AI Business*, 2016, <http://ai.business/2016/06/21/artificial-intelligence-in-defence-and-security-industry/>

288 DARPA, "'Mayhem' Declared Preliminary Winner of Historic Cyber Grand Challenge," 2016, <http://www.darpa.mil/news-events/2016-08-04>; Coldewey, "Carnegie Mellon's Mayhem AI Takes Home \$2 Million from DARPA's Cyber Grand Challenge."

289 Deep Learning Analytics, "Deep Learning Analytics," *Deep Learning Analytics*, 2016, <http://www.deeplearninganalytics.com/single-post/2016/05/15/Deep-Learning-Analytics-Develops-DARPA-Deep-Machine-Learning-Prototype>

290 John Keller, "DARPA TRACE Program Using Advanced Algorithms, Embedded Computing for Radar Target Recognition," *Military & Aerospace Electronics*, July 24, 2015, <http://www.militaryaerospace.com/articles/2015/07/hpec-radar-target-recognition.html>.

291 Ibid.

A more extreme version of these systems could also see the increasing combat competitiveness of UAVs acting as 'wingmen' to human fighter pilots: in the spring of 2016, Psibernetix's artificially intelligent fighter pilot ALPHA soundly defeated US Air Force Colonel (ret.) Gene Lee in a series of simulated dogfights. The 'fuzzy logic' based system was able to process sensor data and plan combat moves in less than a millisecond (more than 250 times faster than the eye can blink), while using no more computing power than that supplied by a \$29 Raspberry Pi.²⁹² Reflecting on the exercise, the experienced air instructor argued that it was "the most aggressive, responsive, dynamic and credible AI I've seen to date."²⁹³

Algorithmic targeting could also increase versatility and accuracy at a theatre level. In the US, the joint targeting cycle currently consists of six major steps – "(1) end state and commander's objectives, (2) target development and prioritizing, (3) capabilities analysis, (4) commander's decision and force assignment, (5) mission planning and force execution, and (6) assessment"²⁹⁴. Each of these steps contains its own feedback mechanisms and time lags. Yet while this approach works well in traditional wars, in complex combat environments, where insurgent networks could reform or quickly adapt to dominant institutional tactics, the amounts of data available on which decisions could be based, can overwhelm traditional analysts, at times actually inhibiting effective, accurate or informed decisionmaking. In this context, integrating machine learning systems to process diverse information can allow for faster, more effective – and more accurate – targeting decisions.

As noted by one analyst, "[t]he computer can identify a targeted building as a hospital instead of an insurgent stronghold by querying a non-governmental organization (NGO) database instead of hoping a commander is aware of that information."²⁹⁵ If applied with care, such systems could make battlefield decisionmaking not only more rapid and responsive in complex or vague environments, but also actually reduce the risk of civilian casualties – although such 'war-algorithms' may again raise their own quandaries in terms of accountability.²⁹⁶

4.4.3. Mission Handoff (A2A; A2D)

Sustaining military efforts for longer periods of time requires regularly rotating units. This may be required by reasons of time (optimizing the operational tempo of the units) or context (changes in the operational environment that require different force bundles;

292 Coby McDonald, "A.I. Downs Expert Human Fighter Pilot In Dogfight Simulation | Popular Science," *Popular Science*, June 27, 2016, <http://www.popsci.com/ai-pilot-beats-air-combat-expert-in-dogfight?src=SOC&dom=tw>. For the study, see also Nicholas Ernest et al., "Genetic Fuzzy Based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions," *Journal of Defense Management* 06, no. 01 (2016), doi:10.4172/2167-0374.1000144

293 M.B. Reilly, "Beyond Video Games: New Artificial Intelligence Beats Tactical Experts in Combat Simulation," *University of Cincinnati Magazine*, June 27, 2016, http://magazine.uc.edu/editors_picks/recent_features/alpha

294 Lewis, "Capturing Flying Insects."

295 Ibid.

296 Cf. Lewis, Blum, and Modirzadeh, "War-Algorithm Accountability."

new equipment entering the force that an outgoing unit lacks; etc.). The United States' Armed Forces define mission handoff as "[T]he process of passing an ongoing mission from one unit to another with no discernible loss of continuity" (e.g. situational awareness, adversary composition, allies, host nation forces, civilian populace, other USG, etc.).²⁹⁷ Mission handoff has proved a major challenge to many Western militaries, especially in high-threat environments where operational security is of paramount importance. Some of the problems include:

- **The short timeframe available for handoff** while in theater, combined with the difficulties of sharing information afterward;
- **The dynamic nature of the operational environment**, creating uncertainty over what information will be relevant;
- **Unwillingness** of combat-weary outgoing units – eager to return home – to spend much time on handoff;
- **Mismatch** in the information or capabilities offered by the outgoing unit, and those needed by the incoming unit;
- **Format mismatch** as Information (e.g. AARs, intel reporting) is often stored on laptops in un- or poorly structured formats.

In the current situation, outgoing units typically meet face-to-face with incoming units and transfer a select number of documents they deem relevant and important. In the short run, AI may offer the incoming unit with an opportunity to have the system ingest all unstructured text generated by the departing unit during its deployment (including After Action Reports, intelligence reports, briefing materials, etc.). Throughout their rotation, they could then query the AI-enhanced knowledge base for insights.

4.4.4. Situational Awareness and Understanding (A2A)

While greatly supporting force projection in asymmetric conflicts, UAVs at present still know a number of operational limitations, such as a low flying speed and vulnerability to air defense systems. Increasing the autonomy of unmanned systems will strengthen their survivability, enable more higher-end performance, and improve their effectiveness at patrolling or monitoring areas. This feeds into an increased ability for militaries or states to cover far greater areas with sensors, at greater cost-effectiveness than human troops. The enhanced situational awareness enabled by more autonomous and survivable drones can strengthen the security of bases and, experts have suggested, could potentially lead to greater stability between states (such as North- and South Korea) by enhancing monitoring of contested areas, reducing the viability of covert or 'hybrid' operations.²⁹⁸

297 United States, Joint Chiefs of Staff, *Foreign Internal Defense*, Joint Publication, JP 3-22 (Washington, D.C.: Joint Chiefs of Staff, 2010), <http://purl.fdlp.gov/GPO/gpo29282>

298 Michael C. Horowitz, Sarah E. Kreps, and Matthew Fuhrmann, "The Consequences of Drone Proliferation: Separating Fact from Fiction," SSRN Scholarly Paper (Rochester, NY: Social Science Research Network, January 25, 2016), <https://papers.ssrn.com/abstract=2722311>.

Conversely, such innovations can also destabilize deterrence. In the spring of 2016, DARPA also trialled the Sea Hunter, the first full-scale prototype of its Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel (ACTUV).²⁹⁹ On the principle that the hardest aspect about countering enemy submarines is not so much tracking them once found, but finding them in the first place, these cost-effective ships are designed to scour the seas for quiet (missile) submarines, and once found automatically trail their quarry globally and for months. The resulting accurate, cost-effective and up-to-date ASW intelligence could have revolutionary – and potentially destabilizing – implications for the (perceived) survivability of the traditionally near-invulnerable seaborne leg of the nuclear triad.³⁰⁰

In a completely different battlefield context, AI systems could greatly increase the safety and efficacy of forces (or indeed non-military government personnel) operating in-the-field, in foreign environments or unfamiliar cultures: systems such as the improved Google Translate or IBM Watson, when combined with natural-language processing and natural voice synthesis capabilities, will be fluent in hundreds of languages, making shortages of interpreters a thing of the past. Moreover, improvements in facial- and emotional recognition systems could improve the ability of soldiers equipped with persistent-stare surveillance modules to gauge a situation, and to ascertain the intent – hostile or peaceful – of many dozens of individuals on crowded streets at a glance.³⁰¹ Applied with care, this could simultaneously support early threat identification, and could reduce the risk of intercultural miscommunications or misperceptions leading to unnecessary or unwitting escalation of interactions with local actors or civilians.

4.4.5. Automated Planning and Manpower Allocation (A2A)

While full-fledged ‘automated planning’ is currently still proving a bottleneck on AI development,³⁰² machine learning systems working from datasets of soldiers’ capacity tests and their past mission performance (individually and in different constellations of teammates) on different types of missions, could formulate elaborate models. On the basis of these, commanders could gain a comprehensive assessment of an individual’s skills, experience, personality, strengths, weaknesses or psychological condition. Machine learning systems could also help them align human talent to mission requirements, and optimize team composition for specific missions, based on expertise & personality, or past unit performance under different conditions. Finally, such systems could draw up force rosters which balance different constraints (e.g. training schedules,

299 Rachel Courtland, “DARPA’s Self-Driving Submarine Hunter Steers Like a Human,” *IEEE Spectrum: Technology, Engineering, and Science News*, April 7, 2016, <http://spectrum.ieee.org/automan/robotics/military-robots/darpa-actuv-self-driving-submarine-hunter-steers-like-a-human>. For project overview, see Scott Littlefield, “Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel (ACTUV),” *DARPA*, accessed September 19, 2016, <http://www.darpa.mil/program/anti-submarine-warfare-continuous-trail-unmanned-vessel>

300 James R. Holmes, “Sea Changes: The Future of Nuclear Deterrence,” *Bulletin of the Atomic Scientists* 72, no. 4 (July 3, 2016): 228–33, doi:10.1080/00963402.2016.1194060.; Cf. *The Economist*, “Anti-Submarine Warfare: Seek, but Shall Ye Find?,” *The Economist*, August 6, 2016, <http://www.economist.com/news/science-and-technology/21703360-proliferation-quieter-submarines-pushing-navies-concoct-better-ways>

301 Cf. Anthony Cruz, “The Robot General: Implications of Watson on Military Operations,” *Armed Forces Journal*, 2011, <http://armedforcesjournal.com/the-robot-general/>

302 Geist, “[Automated] Planning for Tomorrow: Will Artificial Intelligence Get Smarter?”

scheduled downtime, priority missions), ensuring efficient use of service members' time, and indirectly improving morale. hat optimize force time-efficiency and, indirectly, morale.

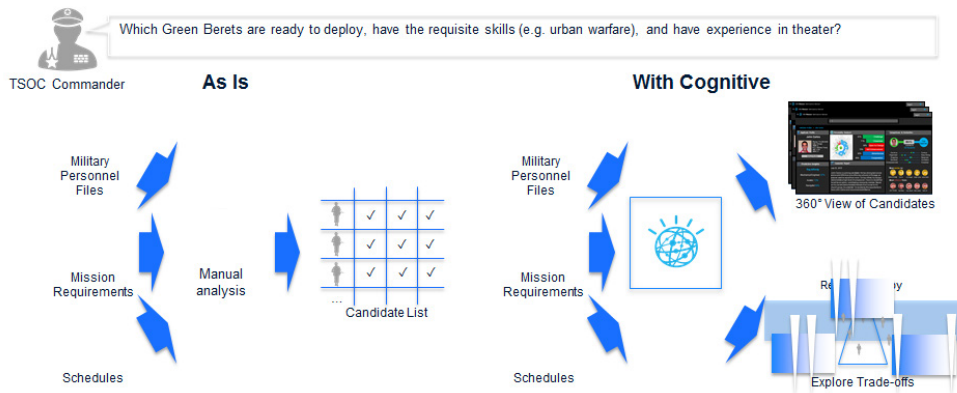


FIGURE 13: AUTOMATED PLANNING AND MANPOWER ALLOCATION


4.4.6. Target Systems Analysis/Target Audience Analysis (A2A, G2G, D2P)

Target Systems Analysis (TSA) and Target Audience Analysis (TAA) are intelligence-related methods used to develop deep understanding of potential areas for operations. They involve the analysis of reports, documents, newsfeeds and other forms unstructured information. Further afield, AI systems could provide probabilistic forecasts of enemy behavior, anticipate and flag bottlenecks or vulnerabilities in supply lines before they occur, and suggest mitigation strategies; draw on data (e.g. weather conditions collected by drones), to examine factors affecting operations and assess the viability of different mission approaches. Natural language processing programs can filter social media and news to identify strategically salient themes³⁰³ – or, conversely, can text-mine the mission reports of own and allied forces in order to identify common themes or patterns in engagements.

For instance, an influential 2015 paper assessed 2,200 military combat incidents involving ISIS, mining these incidents to derive relationships and rules for ISIS vehicle-borne improvised explosive device attacks, as well as indirect fire. By modelling ISIS's behavior, the researchers were able to identify priority targets for the movement, and find strong, previously unrecognized correlations between tactics, such as the fact that spikes in car bombs in Baghdad were often the prelude to ISIS attacks on Northern Iraqi cities, suggesting that ISIS used the bombings as a diversion to draw Iraqi security forces away from prospective targets.³⁰⁴

303 Benjamin Jensen and Ryan Kendall, "Waze for War: How the Army Can Integrate Artificial Intelligence," *War on the Rocks*, September 2, 2016, <https://warontherocks.com/2016/09/waze-for-war-how-the-army-can-integrate-artificial-intelligence/>.

304 Andrew Stanton et al., "Mining for Causal Relationships: A Data-Driven Study of the Islamic State," *arXiv:1508.01192 [Cs]*, August 5, 2015, <http://arxiv.org/abs/1508.01192>. Cf. Alex Lockie, "This Algorithm Could Help Predict ISIS' Future Moves," *Business Insider*, 2015, <http://www.businessinsider.com/machine-learning-used-to>.



Similarly, by drawing on a mix of classified and open-source data, and analyzing these using machine learning algorithms, the CIA in 2016 claimed to have achieved an impressive level of ‘anticipatory intelligence’, being able to anticipate the rise of social unrest and societal instability up to three to five days in advance.³⁰⁵

In this way, the product of the analysis is a sophisticated knowledge of the key individuals and organizations that operate in the economic, cultural, political, tribal, religious spheres of a society. This in turn informs operational planning and deployment, and may also involve the active shaping of public opinion through the careful crafting of messages for specific audiences.

4.4.7. Lessons Learned – Operational and Non-Operational (A2A, A2D)

During and after operations and exercises, defense organizations collect information in order to enable lessons to be developed, shared and learned to help continuous improvement. Intelligent systems (such as IBM’s Watson) could curate repositories and support intelligent query analysis (not unlike many current Google services) enable military personnel to find applicable lessons and apply them to their situation.

Already today, so-called ‘Intelligent Tutoring Systems’ (ITS) are used in the military to develop, cost-effectively deploy responsive, scalable teaching programs that can tailor themselves to the learning needs of each individual.³⁰⁶ For instance, an ITS called SHERLOCK is being used to teach Air Force technicians to run electrical system diagnostics on aircraft,³⁰⁷ and the Information Sciences Institute at the University of Southern California has developed an avatar-based training program to prepare military personnel in appropriate intercultural communication when posted abroad.³⁰⁸

4.4.8. Options Analysis for Capability Development (A2A, A2D)

Options analysis is a key early step in the capability development process. Some governments require that the option set in the initial stages of the acquisition process includes at least one Off-the-Shelf (OTS) solution, where available, as a benchmark. Options that move beyond the requirements of an OTS solution must include a rigorous cost-benefit analysis of the additional capability sought so that the full resource risks and other impacts are understood by governments. As defense must be able to effectively implement any option presented, each option must be achievable in financial, technical, logistics, workforce and schedule terms. The time, effort and expense of examining each

[predict-and-model-isis-2015-9](#)

305 Konkel, “The CIA Says It Can Predict Social Unrest as Early as 3 to 5 Days Out.”

306 Stanford University, “Artificial Intelligence and Life in 2030 One Hundred Year Study on Artificial Intelligence | Report of the 2015 Study Panel.”; p.32.

307 Alan Lesgold et al., “SHERLOCK: A Coached Practice Environment for an Electronics Troubleshooting Job,” in *Computer-Assisted Instruction and Intelligent Tutoring Systems: Shared Goals and Complementary Approaches* (Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1988).

308 Michael V. Yudelson, Kenneth R. Koedinger, and Geoffrey J. Gordon, “Individualized Bayesian Knowledge Tracing Models,” *Artificial Intelligence in Education*, 2013, 171–80.

option in detail makes it essential to concentrate on investigating usually no more than three or four options. Even then that is likely to take 12–18 months to develop the required level of detail. The use of AI may improve the rigour of the analysis and employ trade off analytics to expedite the process

4.4.9. Electronic Medical Records Analysis & Optimizing Medevac (A2A, D2G)

In partnership with the US Veterans Administration IBM's Watson Research team has developed a clinical reasoning prototype called the Electronic Medical Record Analyzer (EMRA). This preliminary technology is designed to ingest a patient's electronic medical record and, using machine-learning techniques, automatically identify and rank the most important health problems for which the patient should be seen. The system is trained to identify symptoms, triggers and risk factors in patient vitals, physician notes, and lab reports. It compares this patient information with the latest medical knowledge in textbooks, medical journals and pharmaceutical information to alert the physician to the patient's most serious health problems.

In a tactical context, narrow AI agents could also aid in the medevac of injured personnel from insecure areas: such systems can query up-to-date databases and cross-reference these against live intelligence from other forces in the area. Combining information on the severity of injuries, the length and security of available exfiltration routes, landing sites and weather conditions, and the projected rate of medical emergencies in the coming days, such a system can perform a preliminary triage and determine the optimal means of evacuating casualties, increasing the efficiency and safety of medical evacuations – and saving lives.³⁰⁹

4.4.10. Documents Classification and 'Crypto-Preserving' Intelligence Sharing (D2G, G2G)

During and after operations and exercises, defense organizations have a need to move information between security domains. This requires the careful checking of the content to ensure no highly classified data or information will be moved or revealed to a less highly classified user or system. While aspects can be automated, the checking often requires a human to conduct the check. This can be very time consuming and is prone to inaccuracy. Watson may be able to more accurately read, understand and verify the content is safe to send from one domain to another, and also help to minimize data aggregation risks – and reduce the risk of inadvertent leaks. Perhaps more importantly, Google has recently conducted research into 'privacy-preserving deep learning',³¹⁰ which could enable multiple organizations (such as hospitals) to share and combine sensitive data (e.g. patient information) on which to train overarching machine learning systems, without actually divulging the underlying information.³¹¹

309 Jensen and Kendall, "Waze for War."

310 Reza Shokri and Vitaly Shmatikov, "Privacy-Preserving Deep Learning" (ACM Press, 2015), 1310–21, doi:10.1145/2810103.2813687

311 Cf. H. Brendan McMahan et al., "Communication-Efficient Learning of Deep Networks from Decentralized

Likewise, using a technique called homomorphic encryption, Microsoft has developed what it calls 'CryptoNets', which are trained deep-learning systems which can take encrypted data and spit out encrypted (but accurate) answers.³¹² The potential applications of such a system for intelligence sharing could be revolutionary – enabling allied intelligence services to enjoy most of the strategic and operational benefits of freely pooling their information and data (e.g. allowing the analysis of global patterns in terrorist communications, tactics or strategy), without the fear that leaks or double agents in their allies' services might compromise that information or its sources.

4.4.11. 'P4' Conflict Prevention – Predictive, Preventive, Personalized and Participatory (X2I)

Critically, whereas most of the previous use cases fall within the more traditional industrial-age 'defense' box the authors of this report suspect that AI's most fundamental (and disruptive) impact may lie elsewhere. The first one we offer resides in the area of prevention. Many defense efforts are currently still predominantly focused on the response stage of the conflict cycle. Armed forces, so goes the argument, are there to wage and win wars or to prevail in conflicts short of war.

Many – also military – voices are increasingly suggesting that we may have to start focusing more on prevention. In many ways, this trend closely mirrors an analogous development in the medical field which is slowly but steadily shifting its focus from punctual interventions to 'fight disease' to more systemic approaches to prevent disease and stimulate healthy lifestyle choices.

One strand of this new thinking has gained some popularity under the acronym 'P4': Predictive, Preventive, Personalized and Participatory medicine'. Big data sets – from molecular and cellular data, over conventional medical data, to vast amounts of new imaging, demographic and environmental data – that are constantly analyzed by machine learning algorithms are expected to allow a radically different, more proactive and 'systemic' approaches to health care that may reverse the ever escalating costs of healthcare and lead to better health outcomes.³¹³

Could we imagine a similar trend in defense and security whereby the combination of big data, ever more powerful computational and AI-algorithmic capabilities leads to deeper insights into the (even individual) drivers of conflict and whereby AI can then also be used to nudge potential agents of conflict away from realizing their malign intentions?

Data," *arXiv:1602.05629* [Cs], February 17, 2016, <http://arxiv.org/abs/1602.05629>

312 Nathan Dowlin et al., "CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy," *Microsoft Research*, February 8, 2016, <https://www.microsoft.com/en-us/research/publication/cryptonets-applying-neural-networks-to-encrypted-data-with-high-throughput-and-accuracy/>; cf. Tom Simonite, "Microsoft and Google Want to Let Artificial Intelligence Loose on Our Most Private Data," *MIT Technology Review*, 2016, <https://www.technologyreview.com/s/601294/microsoft-and-google-want-to-let-artificial-intelligence-loose-on-our-most-private-data/>; Tom Simonite, "A Cloud That Can't Leak," *MIT Technology Review*, 2011, <https://www.technologyreview.com/s/424942/a-cloud-that-cant-leak/>.

313 Walhout, et al., 2013; Flores, et al., 2013; Tian, et al., 2012; Hood, & Friend, 2011

Figure 13 describes the example of Alphabet (Google) Jigsaw's Redirect Method that is already putting this idea in practice, as described by *Wired* magazine.

Google has built a half-trillion-dollar business out of divining what people want based on a few words they type into a search field. In the process, it's stumbled on a powerful tool for getting inside the minds of some of the least understood and most dangerous people on the Internet: potential ISIS recruits. Now one subsidiary of Google is trying not just to understand those would-be jihadis' intentions, but to change them.

Jigsaw, the Google-owned tech incubator and think tank—until recently known as Google Ideas—has been working over the past year to develop a new program it hopes can use a combination of Google's search advertising algorithms and YouTube's video platform to target aspiring ISIS recruits and ultimately dissuade them from joining the group's cult of apocalyptic violence. The program, which Jigsaw calls the Redirect Method and plans to launch in a new phase this month, places advertising alongside results for any keywords and phrases that Jigsaw has determined people attracted to ISIS commonly search for. Those ads link to Arabic- and English-language YouTube channels that pull together preexisting videos Jigsaw believes can effectively undo ISIS's brainwashing—clips like testimonials from former extremists, imams denouncing ISIS's corruption of Islam, and surreptitiously filmed clips inside the group's dysfunctional caliphate in Northern Syria and Iraq.

"This came out of an observation that there's a lot of online demand for ISIS material, but there are also a lot of credible organic voices online debunking their narratives," says Yasmin Green, Jigsaw's head of research and development. "The Redirect Method is at its heart a targeted advertising campaign: Let's take these individuals who are vulnerable to ISIS' recruitment messaging and instead show them information that refutes it."

The results, in a pilot project Jigsaw ran early this year, were surprisingly effective: Over the course of about two months, more than 300,000 people were drawn to the anti-ISIS YouTube channels. Searchers actually clicked on Jigsaw's three or four times more often than a typical ad campaign. Those who clicked spent more than twice as long viewing the most effective playlists than the best estimates of how long people view YouTube as a whole. And this month, along with the London-based startup Moonshot Countering Violent Extremism and the US-based Gen Next Foundation, Jigsaw plans to relaunch the program in a second phase that will focus its method on North American extremists, applying the method to both potential ISIS recruits and violent white supremacists...

But Green says that the Redirect Method, beyond guiding ISIS admirers to its videos, doesn't seek to track them further or identify them, and isn't designed to lead to arrests or surveillance, so much as education. "These are people making

decisions based on partial, bad information,” says Green. “We can affect the problem of foreign fighters joining the Islamic State by arming individuals with more and better information.” She describes the campaign’s work as a kind of extension of Google’s core mission “to make the world’s information accessible and useful.” Perhaps one of world’s most dangerous problems of ignorance and indoctrination can be solved in part by doing what Google does best: Helping people find what they most need to see.³¹⁴”

4.4.12. Empowering Societal Resilience (G2I, E2I)

In previous work,³¹⁵ HCSS has introduced the notion that there are two flip-sides to the security coin: the (highly mediatized) side of the ‘agents of conflict’, and the (surprisingly far less mediatized) side of the ‘agents of (societal) resilience’. We have argued that societal resilience to conflict – the healthy fibers in today’s societies that contain organic antibodies to any excesses of violence – has increased over time and in many senses only continues to strengthen.³¹⁶ We have also submitted that any strategic net assessment of our current security environment and of the available options portfolio that would allow us to achieve our defense and security objectives in a more sustainable way should strive to strike a better balance between the realm of conflict on the one hand (preventing and stopping conflict), and on the other hand the realm of resilience (stimulating both ‘home’ and ‘forward’ resilience).

We want to stress that in our usage of the term ‘resilience’, it applies to both our own societies and to other societies that may be on their way to erupting in violent conflict or to pose a threat to international stability. Resilience is – maybe surprisingly – a relative newcomer in the defence realm. During the Cold War (or before that – e.g.) our own societal resilience did play a significant role in our defense planning efforts. In recent decades it did much less so, although recent security events like terrorism, refugees and populism are once again raising its visibility. As Europe is increasingly concerned about the security effects that radiate inwards from its neighborhood, more (also cost-) effective actionable options to strengthen the security resilience in our neighbors are once again gaining political visibility.

314 Andy Greenberg, “Google’s Clever Plan to Stop Aspiring ISIS Recruits,” *WIRED*, September 7, 2016, <https://www.wired.com/2016/09/googles-clever-plan-stop-aspiring-isis-recruits/>.

315 De Spiegeleire, et al., 2016; De Spiegeleire, & Sweijjs, 2017; De Spiegeleire, et al., 2015

316 Diamandis, & Kotler, 2012; Oosterveld, et al., 2015; Pinker, 2011; Nordberg, 2016

	The Conflict Model	The Resilience Model
Agents	Agents of conflict	Agents of resilience
Goals of main agents	Disruption	Normalcy
'Glue' between agents	Strong	Weak
Organizational form	(Mostly) Organized group	(Maybe networked) Individuals
Efforts by 'us'	(Responsive) Ops/PsyOps	(Preventive) Resilience/nurturing efforts
Action by 'us'	Force	Nudge
('Our') Desired effects	Stop conflict/'win'	Nurture conflict immunity

FIGURE 14: CHARACTERISTICS OF THE CONFLICT MODEL VS THE RESILIENCE MODEL

If we connect this idea of focusing more of our public defense and security efforts on resilience as opposed to conflict to the different layers(/generations) of AI we have described in this think piece, an entirely new options horizon opens up. As we described in the previous use case, our ability to track and understand the dynamics that lead to either strengthened or weakened societal resilience based on a combination of big data, ever more powerful computational and AI-algorithmic capabilities is likely to increase exponentially. Can we imagine our armed forces catalyzing a broader network of sensors and effectors within the defense and security ecosystem that might be able to stimulate both our own and others' 'immune system' against destructive conflict? Would we not want to make sure that our defense (and capability) planning efforts can better adjudicate the claims over scarce resources coming from the more traditional 'conflict- and war-centric' constituencies within our armed forces and defense organizations vs from these more 'resilience-centric' ones?

As we noted in our section about the caveats surrounding AI, there are various legal, ethical and other (even existential) issues that our societies and our polities will have to think through before embarking upon any large-scale efforts to apply AI to 'defense and security'. Many of these caveats apply to both the more traditional defense areas as well as to the prevention and the resilience aspects of 'defense AI'. We would still submit, however, that these last two use cases may prove to be far more palatable to our societies (and polities) than the applications of AI that aim for 'autonomous' lethal strike. The current campaign against 'killer drones' already heralds fundamental and comprehensible societal apprehensions about applying AI to the 'kill-chain'. We suggest that re-focusing that debate on how AI could be used in a more positive way for both prevention and resilience capabilities might prove far more productive.

5 CONCLUSION

5. Conclusion

Artificial Intelligence is highly likely to radically transform our thinking about and our practical approaches towards armed force and defense. The following figure summarizes the main message of this report.

	Artificial Narrow Intelligence	Artificial General Intelligence	Artificial Super Intelligence
'Armed Force' (A2A)	Monitor front-runners and buy opportunistically	Review robustness current force structure	
Ministry of Defense (A2D and D2D)	Identify short-term challenges and opportunities	Identify long-term challenges and opportunities	
Comprehensive (D2G and G2G)	Identify short-term challenges and opportunities	Identify long-term challenges and opportunities	Catalyze dialogue with all stakeholders
Defense and Security Ecosystem (E2E and E2I)	Explore new niches	Explore new niches	Existential challenges and opportunities: fundamental rethink of defense

FIGURE 15: ARTIFICIAL INTELLIGENCE AND 4 LAYERS OF DEFENSE COPY

We have argued in this report that we see 3 different 'types' (and generations) of AI, as well as 4 different 'types' (and generations) of defense and armed force. Figure 15 tries to combine these in one visual. This table in essence represents the option space within which DSOs can choose to prioritize their defense and security AI investments. The remainder of this concluding section will present some thoughts about the advantages and disadvantages of these various options from the point of view of a pro-active but modest security and defense value provider.

Most of the literature we found on AI and defense is currently concentrated in the top left cell of this table: conveniently 'narrow' AI that can help our current operators and warfighters to stay within the current paradigm but to better accomplish their current tasks. We have seen that some of the key military players today – none more than the


United States, but other (aspirational) peer-competitors as well – are focusing primarily on this cell in the table. Based on the recent track record of cost inflation in the A2A option space³¹⁷, the financial implications of focusing on this particular type of defense AI is highly likely to prove exorbitant. Given the Pentagon’s superior financial firepower (which has also just received a significant boost from the incoming Trump administration), its currently still dominant position in the field of AI, and the peculiar dynamics of its political economy of defense – the investments it is likely to make in this area are – in our assessment – likely to be both substantial and impactful.

This means they will yield uniquely new and powerful – even if ‘just’ in the industrial-age sense – capability options that will be expensive but may prove to still provide far better value-for-money than anything European force providers (certainly small- to medium-sized ones) would likely be able to generate in this space in their own right. We therefore recommend a cautiously pragmatic and opportunistic attitude towards this cell of the AI option-space. Investment opportunities will undoubtedly emerge in this segment that small- to medium-sized defense providers might be able to jump on – even with respect to their own R&D funds, but our suggestion here would be to be as opportunistic as our defense industrial basis has been in recent ‘big ticket’ defense procurement projects.

As the overall table and especially the right-downwards orange arrow in [Figure 15](#) suggest, however, we anticipate far superior value-for-money opportunities for SMC DSOs in the other cells of this table – certainly downwards towards the defense and security ecosystem, but arguably even towards the right bottom as we move towards artificial superintelligence.

In first instance, we recommend that small- and medium-sized force providers point their attention and focus downwards in this table. In the second (A2D/D2D) layer the dynamics of the market for solutions seem much less atypical than in the first layer. In the first layer, issues like classification/secretcy and the (connected) *de facto* oligopolistic market structure in many segments of this A2A market lead to highly politicized – and accordingly expensive – capability and acquisition choices. These capabilities are – thankfully – not fungible between the civilian and the military sectors, but that difference – regrettably – comes at a very high price. The second layer, however, is already much more ‘dual-use’, although the current authors never cease to be amazed at the degree to which even this layer seems subject to strikingly similar pathologies as in the first layer (the SAP debacle in many defense organizations being a point in case). We still suspect, however, that many functions that are currently being fulfilled by civil servants in our Ministries of Defense (HR, procurement, inventory management, etc. – but over time also knowledge management,

317 Keith Hartley, “UK Defence Inflation and Cost Escalation,” *Defence and Peace Economics* 27, no. 2 (March 3, 2016): 184–207, doi:10.1080/10242694.2015.1093757; Edward G. Keating and Mark V. Arena, “Defense Inflation: What Has Happened, Why Has It Happened, and What Can Be Done about It?,” *Defence and Peace Economics* 27, no. 2 (March 3, 2016): 176–83, doi:10.1080/10242694.2015.1093760. Commenting on the continually increasing costs of military aircraft, in 1979, Norman Augustine famously quipped “in the year 2054, the entire defense budget will buy just one tactical aircraft. This aircraft will have to be shared between the Air Force and Navy 3-1/2 days each per week”. Norman R. Augustine, “Augustine’s Laws and Major System Development Programs,” *Defense Acquisition Research Journal: A Publication of the Defense Acquisition University* 22, no. 1 (January 2015): 2–6




capability development and management, etc.) may increasingly be encroached upon by (learning) AI algorithms.

If we look at the third layer of [Figure 15](#) – the whole-of-government one – we also anticipate an enormous market for learning algorithms to improve (and rationalize) the coherence and management of the massive amount of paperwork our bureaucracies generate. Many of the real-life connections (or disconnects) between various defense and/or security-relevant policy dossiers currently go unnoticed in our linearly structured administrations. Are European civil servants who are dealing with ‘protecting’ European farmers against non-European agricultural imports cognisant of the direct defense and security impact these policies have the youth cohorts in our North-African neighbor countries? Is the knowledge generated and curated by people working in our educational establishments to make sure our immigrant students acquire the skills required to thrive in modern-day economies and societies even accessible to our development aid workers who try to achieve similar goals in other countries? Might AI help us in overcoming some of the most dysfunctional (yet all too human) aspects of our siloed administrations?

In our own thinking, some of the most promising defense and security AI applications – also (and maybe even especially) for small- and medium-sized DSOs – are likely to emerge in and be increasingly focused on the fourth layer of defense: the defense and security ecosystem layer. We described in chapter 4.2.4 [page 74](#) how in the private sector, the world’s largest companies are in the process of morphing themselves from pipeline businesses into platform businesses surrounded by much broader ecosystems that co-create value for both the platform companies and their ecosystem partners. Artificial intelligence is playing an increasingly powerful role in this through a combination of deep learning, large datasets and high performance computing. We anticipate a similar trend in defense and security. As we pointed out, some important players in this field (like China and Russia) are already embracing/exploiting opportunities in this realm to further what they perceive to be (or construct as) their interest. It is still the case, however, that most of the world’s cutting-edge AI research and start-ups are located in ‘the West’. Even Baidu’s primary AI lab is in Silicon Valley. A more systematic dialogue between our DSOs and various AI ecosystem players may therefore increasingly prove to be a minimal critical requirement for providing unique defense value for defense money.

Alongside moving our focus downwards in [Figure 15](#), this study also recommends moving efforts towards the right. As AI moves towards more AGI-types of comprehensive solutions to the right of [Figure 15](#), the still relatively neat lines between defense and non-defense technologies and applications that we still see today are likely blur even more. This will work both ways. On the one hand, non-defense and security AI-augmented technologies and applications will increasingly have to worry about security and privacy concerns and will have to evolve sustainably reliable solutions for those that might also prove useful for DSOs (and – maybe increasingly – DSEs). And at the same time, more purely defense and security AI-augmented technologies and applications will increasingly be able to piggyback on (and even merge with) the core technologies and applications developed



by their civilian counterparts. To give a concrete example: as medicine starts moving more and more towards what some now call P4 medicine³¹⁸ (personalized, predictive, preventative, participative - see also section 4.4.11 in this report), the sensors, data and algorithms that will be used for health purposes are highly likely to prove equally useful for defense and security purposes.

A continued broader shift in the balance of fiscal power away from defense towards non-defense budget items may very play a major role in this trend. In stark contrast to the Cold War period, where – our DSOs played a critical role in the ebbs and flows of AI as we saw in section 2.2 – our defense and security organizations are quite unlikely to be able to compete with the (current or future) likes of Amazon, Apple, Baidu, Google or Facebook in terms of human, physical or digital capital.³¹⁹ This may even level the playing field for more nimble and forward-looking smaller – and possibly different – DSOs to explore how they might be able to leverage of these new big players to achieve better defense and security solutions in this market segment in ways that some of the bigger states might not.

Obtaining societal buy-in for a whole-of-government role in overseeing and possibly even fusing these AI-generated insights (and actionable policy options) will prove of critical importance. This means that the privacy, legal and ethical implications will have to be properly discussed and addressed in advance in open(-minded) discussions between all of the stakeholders. We suspect that an early debate on more government-led prevention- and resilience-enhancing applications of AI might prove far less acriminious than previous ones on the terrorism threat-related widening of government surveillance competencies or on killer drones. We have to admit, however, that we are not in a position to assess to what extent initial ill-advised initiatives in these afore-mentioned areas may have poisoned the proverbial well.


So far, we have still looked at these three first layers primarily from an ‘us’ vs ‘them’ point of view: us versus them. In our own view, however, the greatest opportunities for AI lie in the bottom E2I layer.

In the wording that we have used throughout this report, we see Google with its Jigsaw example as a great example of a post-industrial agent (in essence one of the most powerful non-state security actors: Alphabet/Google³²⁰) using AI-enhanced post-industrial ‘arms’

318 Medicine will move from a reactive to a proactive discipline over the next decade—a discipline that is predictive, personalized, preventive and participatory (P4). P4 medicine will be fueled by systems approaches to disease, emerging technologies and analytical tools. There will be two major challenges to achieving P4 medicine—technical and societal barriers—and the societal barriers will prove the most challenging.” See Leroy Hood and Stephen H. Friend, “Predictive, Personalized, Preventive, Participatory (P4) Cancer Medicine,” *Nature Reviews Clinical Oncology* 8, no. 3 (2011): 184–187; Mauricio Flores et al., “P4 Medicine: How Systems Medicine Will Transform the Healthcare Sector and Society,” *Personalized Medicine* 10, no. 6 (2013): 565–576; Patricia Sobradillo, Francisco Pozo, and Àlvar Agustí, “P4 Medicine: The Future around the Corner,” *Archivos de Bronconeumología (English Edition)* 47, no. 1 (2011): 35–40.

319 These companies have so far proved fairly reluctant to ‘just’ succumb to their national governments’ attempts to coopt them in their own purposive efforts.

320 We remind our readers that Alphabet’s Android is installed - and is therefore collecting micro-level data - on almost 85% of the world’s 4.61 billion mobile phones in 2016. [statista.com](https://www.statista.com/statistics/274774/forecast-of-mobile-phone-users-worldwide/), “Mobile Phone Users Worldwide 2013-2019,” *Statista*, 2016, <https://www.statista.com/statistics/274774/forecast-of-mobile-phone-users-worldwide/>




(capability bundles) to 'target' individuals with an unadulterated defense and security objective (de-radicalization). And arguably being more effective at it than our DSOs have been able to achieve with their more industrial-age physical and social technologies. Our more traditional military experiences from Libya to Afghanistan have shown disappointing results. But in this example, artificial intelligence appears to have succeeded in nudging some potential ISIS combattants to at least look at videos that are diametrically opposed to their jihadist beliefs. In light of the ethical, legal and even purely political caveats that we have raised with respect to the march of AI – do we as individual human beings and as societies feel comfortable with these powerful private actors being in the driver's seat of these developments? Would we not at least want some public custodians of the public good – whoever these might be – exercise some oversight over these trends? Our own view on this conundrum is that we will have to pursue a more balanced approach towards sustainably effective *and* ethical solutions in this area.

It seems unlikely to us that our current governments, even the most advanced ones – being the industrial-age social technological inventions that they are – will prove able to keep abreast of the unprecedentedly dynamic AI developments in the private (and increasingly also open-source) realms. But it seems equally questionable to us (also based on their track records to date) that these powerful agents of post-industrial innovation, with their continued strong incentives to pursue their own commercial interests, will safeguard our collective and/or individual values and/or interests. We therefore would suggest that just as in the industrial age the governments of nation-states emerged to become the (aspirationally) prudential custodians of the individuals in their realms – so too might new post-industrial social technologies have to emerge to make sure that human individuals can flourish in unencumbered ways.

This, then, would require a more sustained and comprehensive debate between the various public and private stakeholders within the broader defense and security ecosystem about what the optimal way would be to ensure defense and security value for money in this day and age. The suggestion we make in this paper, is that we see few other actors within those ecosystems that have similar perspectives and aptitudes than our defense organizations. Our suggestion is not that they should be the ones to execute the various broader defense and security efforts that might emerge from these broader discussions. We do suggest, however, that they might be better positioned than many others to kickstart and catalyze this discussion.

The final – existentially important – cell in this option space is the one to right of the option space. Artificial super-intelligence – i.e. intelligence that is superior to that of homo sapiens – is likely to pose quite unique challenges to defense and security planners. We have seen in our *longue durée*³²¹ historical overview of the nexus between intelligence and defense how the non-physical, purely cognitive part of 'human force' proved evolutionarily superior to the sometimes physically stronger but cognitively and

321 David Armitage et al., "Le Retour de La Longue Durée: Une Perspective Anglo-Américaine," in *Annales. Histoire, Sciences Sociales*, vol. 70 (Éditions de l'EHESS, 2015), 289–318, <http://www.cairn.info/revue-Annales-2015-2-page-289.htm>.




socially weaker ‘force’ of other species. If that history is a lesson, then the emergence of super-human intelligence is surely a transition that requires extraordinarily careful human consideration. We agree with many other analysts who have looked into this issue that this is an unusually – maybe even unprecedentedly – wicked challenge. Non-ASI AI – both ANI and AGI – is likely to open up many unique opportunities that are likely to enhance our human – individual and collective – experience. Most of us are likely to welcome that part of it. At the same time, however, we also concur with the majority of analysts who anticipate major – maybe even existential – threats when non-human intelligence exceeds its human equivalent. Who will be *homo sapiens*’ custodian, ‘ambassador’ in this debate? Can we imagine a situation where our respective DSOs/DSEs team up together behind the flag of *homo sapiens* in order to address the AI control issue? If so, who will take the initiative for such an effort? Is it more likely to come from a relatively big or a relatively small country?

Our own summary assessment, therefore, is that developed small- to medium-sized force providers would be better served with a capability, partnership, concepts, and policy portfolio that is focused on the lower (and – depending on the observable future trajectory of the trend towards ASI – right) regions of this options space. They might still be able to benefit from the investments their bigger allies are likely to make – for reasons better known to themselves – in the top-left cell of this option space. At the same time, however, we also suggest that especially the other regions of this option space present far superior value for money from a defense and security point of view than the top-left quadrant. Any DSO who can strike up a privileged relationship with actors like Facebook, Google, IBM, Microsoft, etc. is far more likely to lead to be able to start generating the types of security and defense effects our political leaders request than anything our DSOs are currently pursuing.

If we put AI in the same category in which we have put other supposedly game changing military technologies like ‘precision-guided munitions’, ‘stealth’, ‘network-centrism’ or even ‘cyber’ (especially the way many armed forces continue to look at it today), we risk missing the potentially not just game- but paradigm-changing nature of AI. We started this report by reflecting on the quintessential role that (individual and social) *intelligence* (in the broader sense of the term) has played in the evolutionary development of *homo sapiens*. The argument we tried to develop throughout this report is that AI does not just represent another revolution in military affairs (RMA). Most of these RMAs have been primarily about new physical technologies that allowed very similar armed forces to do surprisingly similar things with suspiciously similar tools – still overwhelmingly within the same paradigm. On occasions, they have also been about social technologies: how armed forces reorganized themselves to better ‘beat’ their enemy – but still very much within the same “us vs them” paradigm.

AI may prove to be different. It is not (primarily) about physical technologies, even if these physical technologies may still, based on our survey of the literature, end up being the first part of the ‘defense enterprise’ that will end up being affected. It is also not primarily



about social or even digital technologies, even if we suspect that in the medium term its impact will make itself be felt much more in those domains than in the physical domain. AI is first and foremost about something more fundamental: a quantum leap in raw intelligence – the largely unsung uber-weapon that propelled homo sapiens to the top of the food chain.

We posit that smaller, nimble, forward-looking DSOs have a unique opportunity to start redefining who they are, what they do, and how they do it. AI may very well prove to be one of the key enablers of this (disruptive) transformation. Its algorithmic and (non-defense-specific) essence means that it is easier (though by no means easy) to ‘jump in’ than is the case for many industrial-kinetic technologies. Building a sixth-generation jet fighter is not in the cards for a small- to medium-sized force provider. Applying existing, open-source, deep-learning AI algorithms to the early detection of various defense and/or security pathologies, or to find out which types of purposive actions might lead to the sustainable de-escalation of an explosive intra- or inter-state altercation, or to discover effective anti-inflammatory capabilities that can be deployed by the international community whenever a political entrepreneur in any country tries to whip up his (or her) group or society in a sectarian frenzy – very much is.

This brings us to the following set of recommendations for small- to medium-sized defense and security organizations:

1. ‘Open the windows’ on how we think about future applications of AI to different (also possibly disruptive) futures of defense and security – and do so not from the point of view of our current mindset, structures, capabilities, partnership choices, etc.; but from the point of view of the broader defense and security objectives we want to sustainably achieve;
2. Identify a prioritized list of AI applications that can be rapidly implemented in the framework of existing structures and processes (‘quick wins’). This effort would benefit from an agent of change that can start initiating, experimenting with and stimulating the implementation of (some of) these applications within the defense organization – especially those that focus more on prevention and resilience capabilities.
3. Start a dialogue with various actors across the entire defense and security ecosystem about both promising opportunities and critical challenges in AI (including legal, ethical and privacy aspects) throughout the option space;
4. Start designing a capability, partnership, concept and policy portfolio that is focused on the lower regions of this defense AI options space;
5. Get and remain situationally aware of what is going on in this exceptionally dynamic field. Acknowledge that this cannot be done the ‘old’ triangular (golden or not) way, but requires engaging the entire defense and security ecosystem;
6. Make sure that new emerging and promising options which may not (yet) have

advocates within our defense and security organizations stand a chance against the bigger-ticket legacy items for which there are (powerful) stakeholders;

7. Start thinking about artificial superintelligence and engage with the community that has started thinking about actionable options in that part of the option space (also recognizing that this may open up new avenues for engaging with actors such as China and/or the Russian Federation).

We are encouraged by the fact that the Dutch DSO has already started moving in this direction. The importance of both artificial intelligence and ecosystems has been highlighted in the department's Strategic Knowledge and Innovation Agenda (SKIA) 2016-2020³²² and in its multi-year perspective on the future of the armed forces (*Houvast in een onzekere wereld*)³²³. The concept of the 'defense and security ecosystem' is increasingly being embraced and internalized by the organization, as most vividly illustrated by the 2017 Future Force Conference³²⁴ that was co-organized by the Minister of Defense and the Chief of Defense and was entirely devoted to the operationalization of this concept³²⁵. Building on these first steps in the right direction, we sincerely hope this report contributes to a more fundamental rethink of how AI may affect the future of defense and security and how our DSOs can responsibly harness developments in this field to achieve more sustainable defense and security solutions.

322 Ministry of Defence, The Netherlands and Jeanine A. Hennis-Plasschaert, "Strategische kennis- en innovatieagenda 2016-2020. Vóórbliven in een onzekere wereld [Strategic Knowledge and Innovation Agenda 2016-2020. Staying Ahead in an Uncertain World]" (*The Hague: Ministry of Defence*, October 2016).

323 Ministerie van defensie [Ministry of Defense, The Netherlands], "Houvast in een onzekere wereld Lijnen van ontwikkeling in het meerjarig perspectief voor een duurzaam gereede en snel inzetbare krijgsmacht [A Grip in an Uncertain World. Lines of Development for a Sustainably Ready and Quickly Deployable Armed Force]," rapport (*The Hague: Ministry of Defence*, February 14, 2017), <https://www.rijksoverheid.nl/documenten/rapporten/2017/02/14/meerjarig-perspectief-krijgsmacht-houvast-in-een-onzekere-wereld>.

324 "Future Force Conference 2017," Future Force Conference 2017, 2017, <http://futureforceconference.com/>.


325 Ministry of Defence, The Netherlands, "Final Report Future Force Conference 2017. From Partnerships to Ecosystems: Combining Our Efforts for a More Secure World" (*The Hague*, 2017), http://futureforceconference.com/wp-content/uploads/2017/04/Web_FFC-rapportage.pdf.


6 BIBLIOGRAPHY

6. Bibliography

- 2015 Study Panel. "Artificial Intelligence and Life in 2030." One Hundred Yearly on Artificial Intelligence. Stanford, 2016. https://ai100.stanford.edu/sites/default/files/ai100report10032016fnl_singles.pdf.
- Abadi, Martin, and David G. Andersen. "Learning to Protect Communications with Adversarial Neural Cryptography," 2016. <https://arxiv.org/pdf/1610.06918v1.pdf>.
- Ackerman, Evan. "Can Winograd Schemas Replace Turing Test for Defining Human-Level AI?" *IEEE Spectrum: Technology, Engineering, and Science News*, July 29, 2014. <http://spectrum.ieee.org/automaton/robotics/artificial-intelligence/winograd-schemas-replace-turing-test-for-defining-humanlevel-artificial-intelligence>.
- AI.Business. "Artificial Intelligence in Defence and Security Industry." *AI Business*, 2016. <http://ai.business/2016/06/21/artificial-intelligence-in-defence-and-security-industry/>.
- Alberti, Chris, Dave Orr, and Slav Petrov. "Meet Parsey's Cousins: Syntax for 40 Languages, plus New SyntaxNet Capabilities." *Google Research Blog*. Accessed August 21, 2016. <https://research.googleblog.com/2016/08/meet-parseys-cousins-syntax-for-40.html>.
- Alberts, David S., and Richard E. Hayes. *Power to the Edge: Command, Control in the Information Age*. Information Age Transformation Series. Washington, DC: CCRP Publication Series, 2003.
- Anderson, Berit, and Brett Horvath. "The Rise of the Weaponized AI Propaganda Machine." *Scout*, 2017. <https://scout.ai/story/the-rise-of-the-weaponized-ai-propaganda-machine>.
- Arbesman, Samuel. "Quantifying the Ease of Scientific Discovery." *Scientometrics* 86, no. 2 (February 2011): 245–50. doi:10.1007/s11192-010-0232-6.
- Armitage, David, Jo Guldi, Jérôme Baudry, and others. "Le Retour de La Longue Durée: Une Perspective Anglo-Américaine." In *Annales. Histoire, Sciences Sociales*, 70:289–318. Éditions de l'EHESS, 2015. <http://www.cairn.info/revue-Annales-2015-2-page-289.htm>.
- Armstrong, Stuart, and Kaj Sotala. "How We're Predicting AI—or Failing to." In *Beyond Artificial Intelligence*, 11–29. Springer, 2015. http://link.springer.com/chapter/10.1007/978-3-319-09668-1_2.
- army-technology.com. "Elbit Delivers Command and Staff Trainer to Royal Netherlands Army - Army Technology," November 23, 2012. <http://www.army-technology.com/news/newselbit-delivers-command-staff-trainer-royal-netherlands-army>.

- "Artificial Intelligence - Lockheed Martin." Accessed September 21, 2016. <http://www.lockheedmartin.com/us/atl/research/artificial-intelligence.html>.
- Aubé, Tony. "AI and the End of Truth." *The Startup*, February 14, 2017. <https://medium.com/swlh/ai-and-the-end-of-truth-9a42675de18#.drxlxqv4z>.
- Augustine, Norman R. "Augustine's Laws and Major System Development Programs." *Defense Acquisition Research Journal: A Publication of the Defense Acquisition University* 22, no. 1 (January 2015): 2–6.
- Avenue, Human Rights Watch | 350 Fifth, 34th Floor | New York, and NY 10118-3299 USA | t 1.212.290.4700. "The Growing International Movement Against Killer Robots." *Human Rights Watch*, January 5, 2017. <https://www.hrw.org/news/2017/01/05/growing-international-movement-against-killer-robots>.
- Bank of America Merrill Lynch. "Robot Revolution - Global Robot & AI Primer." Primer. Thematic Investing. Bank of America Merrill Lynch, December 2015. http://www.bofaml.com/content/dam/boamlimages/documents/PDFs/robotics_and_ai_condensed_primer.pdf.
- Barnes, Dan. "Investing in Artificial Intelligence - Is It the Right Time to Buy?" *Raconteur*, December 15, 2015. <http://raconteur.net/technology/investing-in-artificial-intelligence-is-it-the-right-time-to-buy>.
- Barr, A., and Feigenbaum, eds. *The Handbook of Artificial Intelligence*. Vol. 2. Stanford, California & Los Altos, California: HeurisTech Press and William Kaufmann, 1982.
- Barrett, A.M., S.D. Baum, and K.R. Hostetler. "Analyzing and Reducing the Risks of Inadvertent Nuclear War between the United States and Russia." *Science and Global Security* 21, no. 2 (2013): 106–33.
- Battelle, John. *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Reprint. New York: Portfolio, 2006.
- . *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Portfolio, 2006.
- Baum, S.D. "Confronting the Threat of Nuclear Winter." *Futures* 72 (2015): 69–79.
- Beinhocker, Eric. *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics*. Boston Mass.: Harvard Business School Press, 2006.
- Bell, Daniel. *The Coming of Post-Industrial Society: As Venture in Social Forecasting*. Peregrine Books. Harmondsworth: Penguin Books, 1976.
- Bellman, R.E. *An Introduction to Artificial Intelligence: Can Computers Think?* San Francisco: Boyd & Fraser Publishing Company, 1978.
- Bendett, Samuel. "Can Russia's Military Bots Keep Pace?" *RealClearWorld*, August 27, 2015. http://www.realclearworld.com/blog/2015/08/robots_military_russia_hopes.html.
- . "Get Ready, NATO: Russia's New Killer Robots Are Nearly Ready for War." Text. *The*

- 
- National Interest*, March 7, 2017. <http://nationalinterest.org/blog/the-buzz/russias-new-killer-robots-are-nearly-ready-war-19698>.
- Bergen, Mark. "Mercedes-Benz Wants to Beat Google, Uber to Our Driverless Future." *Recode*, November 26, 2015. <http://www.recode.net/2015/11/26/11620962/mercedes-benz-wants-to-beat-google-uber-to-our-driverless-future>.
- Blanchard, Ben. "China Eyes Artificial Intelligence for New Cruise Missiles." *Reuters*, August 19, 2016. <http://www.reuters.com/article/us-china-defence-missiles-idUSKCN10U0EM>.
- Blasko, Dennis J. "'Technology Determines Tactics': The Relationship between Technology and Doctrine in Chinese Military Thinking." *Journal of Strategic Studies* 34, no. 3 (June 1, 2011): 355–81. doi:10.1080/01402390.2011.574979.
- Boeriu, Horatiu. "BMW Sees New Opportunities in Artificial Intelligence." *BMW BLOG*, June 9, 2016. <http://www.bmwblog.com/2016/06/09/bmw-sees-new-opportunities-artificial-intelligence/>.
- Bornstein, Aaron M. "Is Artificial Intelligence Permanently Inscrutable?" *Nautilus*, 2016. <http://nautil.us/issue/40/learning/is-artificial-intelligence-permanently-inscrutable>.
- Bostrom, Nick. "Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards." *Journal of Evolution and Technology* 9, no. 1 (2002).
- . *Superintelligence: Paths, Dangers, Strategies*. OUP Oxford, 2014.
- Bostrom, Nick, and Milan M. Cirkovic. *Global Catastrophic Risks*. 1 edition. Oxford; New York: Oxford University Press, 2011.
- Bostrom, Nick, and Anders Sandberg. "Cognitive Enhancement: Methods, Ethics, Regulatory Challenges." *Science and Engineering Ethics* 15, no. 3 (September 1, 2009): 311–41. doi:10.1007/s11948-009-9142-5.
- Brockman, Greg, and Ilya Sutskever. "Introducing OpenAI." *OpenAI*, 2015. <https://openai.com/blog/introducing-openai/>.
- Brownstein, Barry J., and et al. "Technological Assessment of Future Battlefield Robotic Applications." Proceedings of the Army Conference on Application of AI to Battlefield Information Management. White Oak: US Navy Surface Weapons Center, 1983.
- Bryant, Ross. "Google's AI Becomes First Non-Human to Qualify as a Driver." *Dezeen*, February 12, 2016. <https://www.dezeen.com/2016/02/12/google-self-driving-car-artificial-intelligence-system-recognised-as-driver-usa/>.
- Bryson, Arthur Earl, and Yu-Chi Ho. *Applied Optimal Control: Optimization, Estimation, and Control*. Blaisdell Publishing Company or Xerox College Publishing, 1969.
- Buchanan, Brenda. *Gunpowder, Explosives and the State : A Technological History*. Aldershot England ;; Burlington VT: Ashgate, 2006.
- Burgess, Matt. "Microsoft's AI Can Detect Your Emotions (but Only If You're Angry)." *WIRED*

- 
- UK, 2015. <http://www.wired.co.uk/article/microsoft-predict-emotions-artificial-intelligence>.
- Business Insider. "The 'Google Brain' Is a Real Thing but Very Few People Have Seen It." *Business Insider*, September 2, 2016. <http://finance.yahoo.com/news/google-brain-real-thing-very-002720762.html>.
- "Campaign to Stop Killer Robots." Accessed March 21, 2017. <https://www.stopkillerrobots.org/>.
- Capala, Matthew. "Machine Learning Just Got More Human with Google's Rankbrain." *The Next Web*, September 2, 2016. <http://thenextweb.com/artificial-intelligence/2016/09/02/machine-learning-just-got-more-human-with-googles-rankbrain/>.
- Carter, Ashton. "Keynote Address: The Path to the Innovative Future of Defense." *Center for Strategic and International Studies*, October 28, 2016. https://csis-prod.s3.amazonaws.com/s3fs-public/event/161028_Secretary_Ashton_Carter_Keynote_Address_The_Path_to_the_Innovative_Future_of_Defense.pdf.
- Castells, Manuel. *The Rise of the Network Society*. 2nd ed., With a new pref. The Information Age : Economy, Society, and Culture, v. 1. Chichester, West Sussex ; Malden, MA: Wiley-Blackwell, 2010.
- Castelvecchi, Davide. "Deep Learning Boosts Google Translate Tool." *Nature News*. Accessed February 25, 2017. doi:10.1038/nature.2016.20696.
- Cath, Corinne J. N., Sandra Wachter, Brent Mittelstadt, Mariarosaria Taddeo, and Luciano Floridi. "Artificial Intelligence and the 'Good Society': The US, EU, and UK Approach." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, December 23, 2016. <https://papers.ssrn.com/abstract=2906249>.
- Caygill, Bex. "AI Predicts Outcomes of Human Rights Trials." *UCL*, October 2016. <http://www.ucl.ac.uk/news/news-articles/1016/241016-AI-predicts-outcomes-human-rights-trials>.
- CB Insights. "Deep Interest In AI: New High In Deals To Artificial Intelligence Startups In Q4'15." *CB Insights - Blog*, February 4, 2016. <https://www.cbinsights.com/blog/artificial-intelligence-startup-funding-trends/>.
- . "From Virtual Nurses To Drug Discovery: 106 Artificial Intelligence Startups In Healthcare." *CB Insights - Blog*, February 3, 2017. <https://www.cbinsights.com/blog/artificial-intelligence-startups-healthcare/>.
- Chalmers, David J. "The Singularity: A Philosophical Analysis." *Journal of Consciousness Studies* 17 (2010): 7–65.
- Chandler, Alfred D. "How High Technology Industries Transformed Work and Life Worldwide from the 1880s to the 1990s." *Capitalism and Society* 1, no. 2 (January 7, 2006). doi:10.2202/1932-0213.1004.

- 
- Charniak, E., and D. McDermott. *Introduction to Artificial Intelligence*. Massachusetts: Addison-Wesley Reading, 1985.
- Chase, Kenneth. *Firearms : A Global History to 1700*. Cambridge UK ;;New York NY: Cambridge University Press, 2003.
- Chi-Dong, Lee. "Gov't to Invest 1 Tln Won in Artificial Intelligence." *Yonhap News Agency*, March 17, 2016. <http://english.yonhapnews.co.kr/business/2016/03/17/81/0501000000AEN20160317003751320F.html?bb441f50>.
- Choudary, Sangeet Paul. *Platform Scale: How an Emerging Business Model Helps Startups Build Large Empires with Minimum Investment*. First edition. S.L.: Platform Thinking Labs Pte. Ltd, 2015.
- Clark, Don. "U.S. Agencies Block Technology Exports for Supercomputer in China." *Wall Street Journal*, April 9, 2015, sec. Tech. <http://www.wsj.com/articles/u-s-agencies-block-technology-exports-for-supercomputer-in-china-1428561987>.
- Clark, Jack. "Google Turning Its Lucrative Web Search Over to AI Machines." *Bloomberg.com*, October 26, 2015. <http://www.bloomberg.com/news/articles/2015-10-26/google-turning-its-lucrative-web-search-over-to-ai-machines>.
- Coldewey, Devin. "Carnegie Mellon's Mayhem AI Takes Home \$2 Million from DARPA's Cyber Grand Challenge." *TechCrunch*, 2016. <http://social.techcrunch.com/2016/08/05/carnegie-mellons-mayhem-ai-takes-home-2-million-from-darpas-cyber-grand-challenge/>.
- Committee on Innovations in Computing and Communications, National Research Council. "Chapter 9: Developments in Artificial Intelligence | Funding a Revolution: Government Support for Computing Research." In *Funding a Revolution: Government Support for Computing Research*, 2008. <https://web.archive.org/web/20080112001018/http://www.nap.edu/readingroom/books/far/ch9.html>.
- Courtland, Rachel. "DARPA's Self-Driving Submarine Hunter Steers Like a Human." *IEEE Spectrum: Technology, Engineering, and Science News*, April 7, 2016. <http://spectrum.ieee.org/autotom/robotics/military-robots/darpa-actuv-self-driving-submarine-hunter-steers-like-a-human>.
- Cox, Mark. "Machine-Learning Based Risk Analytics Vendor Gurucul Launches First Channel Program." *ChannelBuzz.ca*, June 1, 2016. <http://www.channelbuzz.ca/2016/06/machine-learning-based-risk-analytics-vendor-gurucul-launches-first-channel-program-17747/>.
- Crevel, Martin van. "Napoleon and the Dawn of Operational Warfare." In *The Evolution of Operational Art: From Napoleon to the Present*, edited by John Andreas Olsen and Martin van Crevel. Oxford: Oxford University Press, 2011.
- Crever, Daniel. *AI: The Tumultuous History of the Search for Artificial Intelligence*. New York, NY: Basic Books, 1993.

- . *AI: The Tumultuous Search for Artificial Intelligence*. New York: BasicBooks, 1993.
- Cross, Stephen E., and Edward Walker. "DART: Applying Knowledge Based Planning and Scheduling to CRISIS Action Planning." In *Intelligent Scheduling*, edited by Monte Zweben and Mark Fox, 711–29. San Francisco, CA: Morgan Kaufmann, 1994.
- Cruz, Anthony. "The Robot General: Implications of Watson on Military Operations." *Armed Forces Journal*, 2011. <http://armedforcesjournal.com/the-robot-general/>.
- DARPA. "Grant Opportunity - DARPA-BAA-16-53 - Explainable Artificial Intelligence (Xai)." Department of Defense - DARPA - Information Innovation Office, n.d. <http://www.grants.gov/web/grants/view-opportunity.html?oppld=287284>.
- . "'Mayhem' Declared Preliminary Winner of Historic Cyber Grand Challenge," 2016. <http://www.darpa.mil/news-events/2016-08-04>.
- Davis, Ernest. "The Singularity and the State of the Art in Artificial Intelligence." *Ubiquity* 2014, no. October (November 4, 2014): 1–12. doi:10.1145/2667640.
- Dawson, Ashley. *Extinction: A Radical History*. New York: Or Books Llc, 2016.
- De Spiegeleire, Stephan De, and Tim Sweijts. *Volatility and Friction in the Age of Disintermediation. HCSS StratMon 2016-2017 Yearly Report*, 2017.
- De Spiegeleire, Stephan, and Peter Essens. "C2 That! Command and Control over Post-Industrial Armed Forces." Santa Monica CA, 2010.
- . "C2 That! Command and Control over Post-Industrial Armed Forces." In *The Evolution of C2*. Santa Monica CA, 2010.
- De Spiegeleire, Stephan, Willem Th. Oosterveld, Ridder, Marjolein, Tim Sweijts, Frank Bekkers, Polackova, Dana, Ward, Scott, Eldin Salah, Kamal, Rutten, Rik, and Olah, Nathalie. *Si Vis Pacem, Para Utique Pacem. Individual Empowerment, Societal Resilience and the Armed Forces*. The Hague, The Netherlands: The Hague Centre for Strategic Studies, 2015. <http://www.literatuurplein.nl/boekdetail.jsp?boekId=1078494>.
- "Deep Learning Analytics." *Deep Learning Analytics*, 2016. <http://www.deeplearninganalytics.com/single-post/2016/05/15/Deep-Learning-Analytics-Develops-DARPA-Deep-Machine-Learning-Prototype>.
- Defense Science Board. "Defense Science Board - Charter," May 6, 2016. <http://www.acq.osd.mil/dsb/charter.htm>.
- . "Defense Science Board Summer Study on Autonomy," 2016. <https://www.hsd.org/?abstract&did=794641>.
- . "Report of the Defense Science Board Summer Study on Autonomy." Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, June 2016. <https://www.hsd.org/?view&did=794641>.
- Dehaene, Stanislas. *Consciousness and the Brain: Deciphering How the Brain Codes Our*

- Thoughts*. New York, New York: Viking, 2014.
- . *Reading in the Brain: The Science and Evolution of a Human Invention*. New York: Viking, 2009.
- Department for Culture, Media & Sport. “£17 Million Boost for the UK’s Booming Artificial Intelligence Sector - GOV.UK.” Press Release. *Gov.uk*, February 26, 2017. <https://www.gov.uk/government/news/17-million-boost-for-the-uks-booming-artificial-intelligence-sector>.
- . “UK Digital Strategy.” *Gov.uk*, March 1, 2017. <https://www.gov.uk/government/publications/uk-digital-strategy>.
- Department of Defense. “Unmanned Systems Integrated Roadmap: FY2013-2038,” 2013. <http://archive.defense.gov/pubs/DOD-USRM-2013.pdf>.
- Devaine, Marie, Guillaume Hollard, and Jean Daunizeau. “Theory of Mind: Did Evolution Fool Us?” *PLOS ONE* 9, no. 2 (February 5, 2014): e87619. doi:10.1371/journal.pone.0087619.
- DiBona, Chris. “The Limits of Biological Intelligence.” In *What to Think about Machines That Think: Today’s Leading Thinkers on the Age of Machine Intelligence*, edited by John Brockman. HarperCollins, 2015.
- Dobbing, Mary, and Chris Cole. “Israel and the Drone Wars: Examining Israel’s Production, Use and Proliferation of UAVs.” Drone Wars UK, January 2014. <https://dronewarsuk.files.wordpress.com/2014/01/israel-and-the-drone-wars.pdf>.
- Domingos, Pedro. The Knowledge Project: Pedro Domingos on Artificial Intelligence. Interview by Shane Parrish. Podcast, August 30, 2016. <http://theknowledgeproject.libsyn.com/pedro-domingos-on-artificial-intelligence>.
- . *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. 1 edition. New York: Basic Books, 2015.
- . “The Race for the Master Algorithm Has Begun.” *WIRED UK*, 2016. <http://www.wired.co.uk/article/master-algorithm-pedro-domingos>.
- Dowlin, Nathan, Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. “CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy.” *Microsoft Research*, February 8, 2016. <https://www.microsoft.com/en-us/research/publication/cryptonets-applying-neural-networks-to-encrypted-data-with-high-throughput-and-accuracy/>.
- Dunsworth, Holly M., Anna G. Warrener, Terrence Deacon, Peter T. Ellison, and Herman Pontzer. “Metabolic Hypothesis for Human Altriciality.” *Proceedings of the National Academy of Sciences* 109, no. 38 (September 18, 2012): 15212–16. doi:10.1073/pnas.1205282109.
- Ernest, Nicholas, David Carroll, Corey Schumacher, Matthew Clark, Kelly Cohen, and Gene Lee. “Genetic Fuzzy Based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions.” *Journal of Defense Management* 06, no. 01


- (2016). doi:10.4172/2167-0374.1000144.
- Esmailzadeh, Hadi, Emily Blem, Renée St. Amant, Karthikeyan Sankaralingam, and Doug Burger. "Power Limitations and Dark Silicon Challenge the Future of Multicore." *ACM Transactions on Computer Systems* 30, no. 3 (2012). http://www.cc.gatech.edu/~hadi/doc/paper/2012-tocs-dark_silicon.pdf.
- Evans, David S., and Richard Schmalensee. *Matchmakers: The New Economics of Platform Businesses: How One of the Oldest Business Models on Earth Powers the Most Incredible Companies in the World*. Boston, Massachusetts: Harvard Business Review Press, 2016.
- Extance, Andy. "Robot Scientist Discovers Potential Malaria Drug." *Scientific American*, February 2015. <https://www.scientificamerican.com/article/robot-scientist-discovers-potential-malaria-drug/>.
- Faggella, Daniel. "Artificial Intelligence Industry – An Overview by Segment." *TechEmergence.com*, July 25, 2016. <http://techemergence.com/artificial-intelligence-industry-an-overview-by-segment/>.
- . "How Investors Feel About Artificial Intelligence – from 29 AI Founders and Executives." *TechEmergence.com*, August 1, 2016. <http://techemergence.com/how-investors-feel-about-artificial-intelligence-from-29-ai-founders-and-executives/>.
- . "Valuing the Artificial Intelligence Market, Graphs and Predictions for 2016 and Beyond." *TechEmergence.com*, March 7, 2016. <http://techemergence.com/valuing-the-artificial-intelligence-market-2016-and-beyond/>.
- . "Venture Investments in Artificial Intelligence – Trends in 2016 and Beyond – TechEmergence.com." *TechEmergence.com*, November 29, 2016. <http://techemergence.com/investments-in-artificial-intelligence/>.
- . "What Is Artificial Intelligence? An Informed Definition -." *TechEmergence.com*, October 10, 2016. <http://techemergence.com/what-is-artificial-intelligence/>.
- Farfadi, Sachin Sudhakar, Mohammad Saberian, and Li-Jia Li. "Multi-View Face Detection Using Deep Convolutional Neural Networks." *arXiv:1502.02766 [Cs]*, February 9, 2015. <http://arxiv.org/abs/1502.02766>.
- Farquhar, Sebastian, John Halstead, Owen Cotton-Barratt, Stefan Schubert, Haydn Belfield, and Andrew Snyder-Beattie. "Existential Risk: Diplomacy and Governance." Global Priorities Project, 2017. <https://www.fhi.ox.ac.uk/wp-content/uploads/Existential-Risks-2017-01-23.pdf>.
- Felner, Ariel. "The Israeli AI Community." *AI Magazine* 37, no. 3 (October 7, 2016): 118–22.
- Fiott, Daniel. "Europe and the Pentagon's Third Offset Strategy." *The RUSI Journal* 161, no. 1 (January 2, 2016): 26–31. doi:10.1080/03071847.2016.1152118.
- Flores, Mauricio, Gustavo Glusman, Kristin Brogaard, Nathan D. Price, and Leroy Hood. "P4 Medicine: How Systems Medicine Will Transform the Healthcare Sector and Society."

- Personalized Medicine* 10, no. 6 (2013): 565–576.
- Forrest, Connor. “How Distil Networks Uses Machine Learning to Hunt down ‘Bad Bots’ - TechRepublic.” *TechRepublic*, August 3, 2016. <http://www.techrepublic.com/article/how-distil-networks-uses-machine-learning-to-hunts-down-bad-bots/>.
- Forrest, Stephanie, and Melanie Mitchell. “Adaptive Computation: The Multidisciplinary Legacy of John H. Holland.” *Communications of the ACM* 59, no. 8 (August 2016). <http://web.cecs.pdx.edu/~mm/jhh-cacm.pdf>.
- Freedberg, Sydney J. “Centaur Army: Bob Work, Robotics, & The Third Offset Strategy.” *Breaking Defense*, November 9, 2015. <http://breakingdefense.com/2015/11/centaur-army-bob-work-robotics-the-third-offset-strategy/>.
- Freeman, Chris, and Francisco Louçã. *As Time Goes by: From the Industrial Revolutions to the Information Revolution*. Oxford: Oxford University Press, 2002.
- Friedman, Jeffrey A. *The 2006 Lebanon Campaign and the Future of Warfare: Implications for Army and Defense Policy*. Lulu. com, 2011. <https://books.google.com.au/>
- Fukuyama, Francis. *Political Order and Political Decay: From the Industrial Revolution to the Globalization of Democracy*. Farrar, Straus and Giroux, 2014.
- . *The Origins of Political Order: From Prehuman Times to the French Revolution*. Farrar, Straus and Giroux, 2011.
- Fuller, Samuel H., and Lynette I. Millett, eds. *The Future of Computing Performance: Game Over or Next Level?* Washington, D.C: National Academies Press, 2011.
- Furman, Jason. “Is This Time Different? The Opportunities and Challenges of Artificial Intelligence.” presented at the AI Now: The Social and Economic Implications of Artificial Intelligence Technologies in the Near Term, New York, July 7, 2016. https://www.whitehouse.gov/sites/default/files/page/files/20160707_cea_ai_furman.pdf.
- Furness, Dyllan. “AI in Agriculture? Algorithms Help Farmers Spot Crop Disease like Experts.” *Digital Trends*, October 6, 2016. <http://www.digitaltrends.com/computing/ai-crop-disease/>.
- “Future Force Conference 2017.” Future Force Conference 2017, 2017. <http://futureforceconference.com/>.
- Future of Life Institute. “Asilomar AI Principles.” *Future of Life Institute*, 2017. <https://futureoflife.org/ai-principles/>.
- Futurism, Dom Galeon & Christianna Reedy. “A Google Exec Just Claimed the Singularity Will Happen by 2029.” *ScienceAlert*, 2017. <http://www.sciencealert.com/google-s-director-of-engineering-claims-that-the-singularity-will-happen-by-2029>.
- Garland, Alex. *Ex Machina*. Drama, Mystery, Sci-Fi, 2015.

- Gat, Azar. *War in Human Civilization*. OUP Oxford, 2008.
- . *War in Human Civilization*. Oxford University Press, 2008. <http://www.oupcanada.com/catalog/9780199236633.html>.
- Geist, Edward Moore. “[Automated] Planning for Tomorrow: Will Artificial Intelligence Get Smarter?” *Bulletin of the Atomic Scientists* 70, no. 2 (2017): 80–85.
- Gibney, Elizabeth. “Google’s AI Reasons Its Way around the London Underground.” *Nature News*, 2016. doi:10.1038/nature.2016.20784.
- Goertzel, Ben, and Hugo de Garis. “Seeking the Sputnik of AGI - H+ Media.” *H+ Media*, March 30, 2011. <http://hplusmagazine.com/2011/03/30/seeking-the-sputnik-of-agi/>.
- Goertzel, Ben, and Joel Pitt. “Nine Ways to Bias Open-Source AGI Toward Friendliness.” *Journal of Evolution and Technology* 22, no. 1 (2012): 116–31.
- Google. “Algorithms – Inside Search – Google.” Accessed August 21, 2016. <https://www.google.com/insidesearch/howsearchworks/algorithms.html>.
- Google DeepMind. “Health | Google DeepMind.” *Google DeepMind*. Accessed August 21, 2016. <http://deepmind.com/health>.
- Grant, Adam. *Originals: How Non-Conformists Move the World*. Penguin, 2016.
- Greenberg, Andy. “Google’s Clever Plan to Stop Aspiring ISIS Recruits.” *WIRED*, September 7, 2016. <https://www.wired.com/2016/09/googles-clever-plan-stop-aspiring-isis-recruits/>.
- Greimel, Hans. “What’s behind Toyota’s Big Bet on Artificial Intelligence?” *Automotive News*, 2015. <http://www.autonews.com/article/20151109/OEM06/311099937/whats-behind-toyotas-big-bet-on-artificial-intelligence%3F>.
- Gunning, David. “Explainable Artificial Intelligence (XAI).” *DARPA*. Accessed September 23, 2016. <http://www.darpa.mil/program/explainable-artificial-intelligence>.
- HAL 90210. “This Is What Happens When an AI-Written Screenplay Is Made into a Film.” *The Guardian*, June 10, 2016, sec. Technology. <https://www.theguardian.com/technology/2016/jun/10/artificial-intelligence-screenplay-sunspring-silicon-valley-thomas-middleditch-ai>.
- Hamel, Gary, and Michele Zanini. “Excess Management Is Costing the U.S. \$3 Trillion Per Year.” *Harvard Business Review*, September 5, 2016. <https://hbr.org/2016/09/excess-management-is-costing-the-us-3-trillion-per-year>.
- Hanson, Robin. “Good CS Expert Says: Most Firms That Thinks They Want Advanced AI/ML Really Just Need Linear Regression on Cleaned-up Data.” Microblog. *Twitter*, November 28, 2016. <https://twitter.com/robinhanson/status/803302419249725440>.
- . *The Age of Em: Work, Love, and Life When Robots Rule the Earth*. Oxford, New York: Oxford University Press, 2016.

- . “This AI Boom Will Also Bust.” *Overcoming Bias*, 2016. <http://www.overcomingbias.com/2016/12/this-ai-boom-will-also-bust.html>.
- Harari, Yuval N. *Sapiens: A Brief History of Humankind*. First U.S. edition. New York: Harper, 2015.
- . “Why Humans Run the World.” *Ideas.ted.com*, June 16, 2015. <http://ideas.ted.com/why-humans-run-the-world/>.
- Harari, Yuval Noah. *Homo Deus: A Brief History of Tomorrow*. Random House, 2016.
- . *Sapiens: A Brief History of Humankind*. Random House, 2014.
- Harper, Douglas. “Online Etymology Dictionary: Arm.” Accessed September 20, 2016. www.etymonline.com/index.php?term=arm.
- . “Online Etymology Dictionary: Intelligence.” Accessed September 28, 2016. http://www.etymonline.com/index.php?allowed_in_frame=0&search=intelligence.
- Harré, Michael. “Social Network Size Linked to Brain Size.” *Scientific American*, August 7, 2012. <http://www.scientificamerican.com/article/social-network-size-linked-brain-size/>.
- Hartley, Jean. “Innovation in Governance and Public Services: Past and Present.” *Public Money and Management* 25, no. 1 (2005): 27–34.
- Hartley, Keith. “UK Defence Inflation and Cost Escalation.” *Defence and Peace Economics* 27, no. 2 (March 3, 2016): 184–207. doi:10.1080/10242694.2015.1093757.
- Haugeland, J., ed. *Artificial Intelligence: The Very Idea*. Cambridge, Mass: MIT Press, 1985.
- Hedberg, Sara Reese. “DART: Revolutionizing Logistics Planning.” *IEEE Intelligent Systems* 17, no. 3 (May 2002): 81–83. doi:10.1109/MIS.2002.1005635.
- Hermann, Karl Moritz, Tomáš Kořinský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. “Teaching Machines to Read and Comprehend.” *arXiv:1506.03340 [Cs]*, June 10, 2015. <http://arxiv.org/abs/1506.03340>.
- Hinton, G. E., and R. R. Salakhutdinov. “Reducing the Dimensionality of Data with Neural Networks.” *Science* 313, no. 5786 (July 28, 2006): 504–7. doi:10.1126/science.1127647.
- Hoffman, David E. *The Dead Hand: The Untold Story of the Cold War Arms Race and Its Dangerous Legacy*. New York: Anchor Books, 2009. <http://search.ebscohost.com/login.aspx?direct=true&scope=site&db=nlebk&db=nlabk&AN=735904>.
- Holman, Mark. “Artificial Intelligence and the Future: Promise or Peril.” presented at the Silicon Valley Thought Leader Series, June 2, 2016. https://www.atkearney.com/documents/7789801/0/BB_54028_20160212+SVTLS+AI_FINAL.pdf/94492e02-e5c3-4ffb-9e3a-47c9ac8bf93b.
- Holmes, James R. “Sea Changes: The Future of Nuclear Deterrence.” *Bulletin of the Atomic*

- Scientists* 72, no. 4 (July 3, 2016): 228–33. doi:10.1080/00963402.2016.1194060.
- Hood, Leroy, and Stephen H. Friend. “Predictive, Personalized, Preventive, Participatory (P4) Cancer Medicine.” *Nature Reviews Clinical Oncology* 8, no. 3 (2011): 184–187.
- Horowitz, Michael C., Sarah E. Kreps, and Matthew Fuhrmann. “The Consequences of Drone Proliferation: Separating Fact from Fiction.” SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, January 25, 2016. <https://papers.ssrn.com/abstract=2722311>.
- Hsu, Jeremy. “Deep Learning Makes Driverless Cars Better at Spotting Pedestrians.” *IEEE Spectrum: Technology, Engineering, and Science News*, February 9, 2016. <http://spectrum.ieee.org/cars-that-think/transportation/advanced-cars/deep-learning-makes-driverless-cars-better-at-spotting-pedestrians>.
- Hullinger, Jessica. “7 Absolutely Insane Animal Defense Mechanisms.” *Mental Floss*, October 6, 2013. <http://mentalfloss.com/article/12258/7-absolutely-insane-animal-defense-mechanisms>.
- Hutchins, John. “The History of Machine Translation in a Nutshell,” 2005. <http://www.hutchinsweb.me.uk/Nutshell-2005.pdf>.
- Hutter, Marcus. “Can Intelligence Explode?” *arXiv:1202.6177 [Physics]*, February 28, 2012. <https://arxiv.org/pdf/1202.6177.pdf>.
- IBM Research. “IBM Research: A New Kind of Dementia Treatment.” Accessed September 20, 2016. <https://www.research.ibm.com/articles/dementia-treatment-diagnosis.shtml>.
- IEEE. “Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems.” IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems. Institute of Electrical and Electronics Engineers (IEEE), 2017. http://standards.ieee.org/develop/indconn/ec/ead_v1.pdf.
- Isaac, Mike, and Neal E. Boudette. “Ford to Invest \$1 Billion in Artificial Intelligence Start-Up.” *The New York Times*, February 10, 2017. <https://www.nytimes.com/2017/02/10/technology/ford-invests-billion-artificial-intelligence.html>.
- Ismail, Salim, and Michael S. Malone. *Exponential Organizations: Why New Organizations Are Ten Times Better, Faster, and Cheaper than Yours (and What to Do about It)*. 1. ed. A Singularity University Book. New York, NY: Diversion Books, 2014.
- Israel Today. “Israel First to Deploy Autonomous Military Robots.” *Israel Today*, August 25, 2016. <http://www.israeltoday.co.il/NewsItem/tabid/178/nid/29924/Default.aspx>.
- Jane’s. “Aeronautics Control Systems.” *Jane’s Unmanned Aerial Vehicles and Targets*, 22-Mar-17. <http://janes.ihs.com.proxy1.athensams.net/Janes/Display/1318450>.
- . “Elbit Systems Ltd - Command and Staff Training Systems.” *Jane’s Simulation and Training Systems*, March 23, 2017. <http://janes.ihs.com.proxy1.athensams.net/Janes/Display/1592554>.

- 
- Jensen, Benjamin, and Ryan Kendall. "Waze for War: How the Army Can Integrate Artificial Intelligence." *War on the Rocks*, September 2, 2016. <https://warontherocks.com/2016/09/waze-for-war-how-the-army-can-integrate-artificial-intelligence/>.
- Jesus, Cecille de. "AI Lawyer 'Ross' Has Been Hired By Its First Official Law Firm." *Futurism*, May 11, 2016. <https://futurism.com/artificially-intelligent-lawyer-ross-hired-first-official-law-firm/>.
- Jing, Meng. "China's First 'deep Learning Lab' Intensifies Challenge to US in Artificial Intelligence Race." *South China Morning Post*, February 21, 2017. <http://www.scmp.com/tech/china-tech/article/2072692/chinas-first-deep-learning-lab-intensifies-challenge-us-artificial>.
- Johnson, Theodore R. "Donald Trump's Pentagon and the Future of the Third Offset Strategy: Will the Department of Defense Invest in People or Technology?" *The Atlantic*, November 29, 2016. <https://www.theatlantic.com/politics/archive/2016/11/trump-military-third-offset-strategy/508964/>.
- Juskalian, Russ. "Practical Quantum Computers." *MIT Technology Review* 120, no. 2 (April 3, 2017): 76–81.
- Kania, Elsa. "China May Soon Surpass America on the Artificial Intelligence Battlefield." Text. *The National Interest*, February 21, 2017. <http://nationalinterest.org/feature/china-may-soon-surpass-america-the-artificial-intelligence-19524>.
- . "The Next U.S.-China Arms Race: Artificial Intelligence?" *The National Interest*, 2017. <http://nationalinterest.org/feature/the-next-us-china-arms-race-artificial-intelligence-19729>.
- Kapás, Judit. "Industrial Revolutions and the Evolution of the Firm's Organization: An Historical Perspective." *Journal of Innovation Economics & Management*, no. 2 (2008): 15–33.
- Keating, Edward G., and Mark V. Arena. "Defense Inflation: What Has Happened, Why Has It Happened, and What Can Be Done about It?" *Defence and Peace Economics* 27, no. 2 (March 3, 2016): 176–83. doi:10.1080/10242694.2015.1093760.
- Keller, John. "DARPA TRACE Program Using Advanced Algorithms, Embedded Computing for Radar Target Recognition." *Military & Aerospace Electronics*, July 24, 2015. <http://www.militaryaerospace.com/articles/2015/07/hpec-radar-target-recognition.html>.
- Kelly, J. E. "Computing, Cognition and the Future of Knowing." *Whitepaper, IBM Research*, 2015. http://www.academia.edu/download/44915350/Computing_Cognition_WhitePaper.pdf.
- Kirchner, Lauren. "When Discrimination Is Baked Into Algorithms." *The Atlantic*, September 6, 2015. <https://www.theatlantic.com/business/archive/2015/09/discrimination-algorithms-disparate-impact/403969/>.
- Kleman, Alexandra. "Cooking with Chef Watson, I.B.M.'s Artificial-Intelligence App." *The New Yorker*, 2016. <http://www.newyorker.com/magazine/2016/11/28/cooking-with>

- chef-watson-ibms-artificial-intelligence-app.
- Knight, Will. "Baidu System Rivals People at Speech Recognition." *MIT Technology Review*, 2015. <https://www.technologyreview.com/s/544651/baidus-deep-learning-system-rivals-people-at-speech-recognition/>.
- Kolbert, Elizabeth. *The Sixth Extinction: An Unnatural History*. First edition. New York: Henry Holt and Company, 2014.
- König, Alexandra, Aharon Satt, Alexander Sorin, Ron Hoory, Orith Toledo-Ronen, Alexandre Derreumaux, Valeria Manera, et al. "Automatic Speech Analysis for the Assessment of Patients with Predementia and Alzheimer's Disease." *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring* 1, no. 1 (March 2015): 112–24. doi:10.1016/j.dadm.2014.11.012.
- Konkel, Frank. "The CIA Says It Can Predict Social Unrest as Early as 3 to 5 Days Out." *Defense One*, 2016. <http://www.defenseone.com/technology/2016/10/cia-says-it-can-predict-social-unrest-early-3-5-days-out/132121/>.
- Koza, John R. "Survey of Genetic Algorithms and Genetic Programming." In *WESCON/95. Conference record. Microelectronics Communications Technology Producing Quality Products Mobile and Portable Power Emerging Technologies*, 589. IEEE, 1995. <http://ieeexplore.ieee.org/abstract/document/485447/>.
- Kurzweil, R. *The Age of Intelligent Machines*. Cambridge, Mass: MIT Press, 1990.
- Kurzweil, Ray. *How to Create a Mind: The Secret of Human Thought Revealed*. 1st edition. New York: Viking, 2012.
- . "Long Live AI." *Forbes.com*, August 15, 2015. http://www.forbes.com/home/free_forbes/2005/0815/030.html.
- . *The Singularity Is Near*. Viking Press, 2005.
- Lappin, Yaakov. "Interview: Gadi Shamni, Head of Land Systems, Israel Aerospace Industries." *Jane's Defence Weekly*, September 7, 2016. <http://janes.ihs.com.proxy1.athensams.net/Janes/Display/1782846>.
- Legg, Shane, and Marcus Hutter. "A Collection of Definitions of Intelligence." *arXiv:0706.3639 [Cs]*, June 25, 2007. <http://arxiv.org/abs/0706.3639>.
- . "A Formal Definition of Intelligence for Artificial Systems." *Proc. 50th Anniversary Summit of Artificial Intelligence*, 2006, 197–198.
- . "Universal Intelligence: A Definition of Machine Intelligence." *Minds and Machines* 17, no. 4 (2007): 391–444.
- Lei, Zhao. "Nation's next Generation of Missiles to Be Highly Flexible." *China Daily*, 2016. http://www.chinadaily.com.cn/china/2016-08/19/content_26530461.htm.
- Lenat, Douglas, and R.V. Guha. *Building Large Knowledge-Based Systems*. Addison-Wesley,

- 1989.
- Lesgold, Alan, Suzanne Lajoie, Marilyn Bunzo, and Gary Eggan. "SHERLOCK: A Coached Practice Environment for an Electronics Troubleshooting Job." In *Computer-Assisted Instruction and Intelligent Tutoring Systems: Shared Goals and Complementary Approaches*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1988.
- Levy, Steven. *In the Plex: How Google Thinks, Works, and Shapes Our Lives*. 1st Simon & Schuster hardcover ed. New York: Simon & Schuster, 2011.
- Lewis, Charlie. "Capturing Flying Insects: A Machine Learning Approach to Targeting." *War on the Rocks*, September 6, 2016. <http://warontherocks.com/2016/09/capturing-flying-insects-a-machine-learning-approach-to-targeting/>.
- Lewis, Dustin A., Gabriella Blum, and Naz K. Modirzadeh. "War-Algorithm Accountability." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, August 31, 2016. <http://papers.ssrn.com/abstract=2832734>.
- Lewis-kraus, Gideon. "The Great A.I. Awakening." *The New York Times*, December 14, 2016. <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html>.
- Lighthill, James. "Artificial Intelligence: A General Survey." Artificial Intelligence: A Paper Symposium. Science Research Council, 1973. <http://www.math.snu.ac.kr/~hicho/infomath/Articles/Lighthill%20Report.pdf>.
- Linn, Allison. "Microsoft Researchers Win ImageNet Computer Vision Challenge." *Next at Microsoft*, December 10, 2015. <https://blogs.microsoft.com/next/2015/12/10/microsoft-researchers-win-imagenet-computer-vision-challenge/>.
- Littlefield, Scott. "Anti-Submarine Warfare (ASW) Continuous Trail Unmanned Vessel (ACTUV)." *DARPA*. Accessed September 19, 2016. <http://www.darpa.mil/program/anti-submarine-warfare-continuous-trail-unmanned-vessel>.
- Lockheed Martin. "LRASM: Overview." *Lockheed Martin*, 2016. <http://www.lockheedmartin.com/us/products/LRASM/overview.html>.
- Lockie, Alex. "This Algorithm Could Help Predict ISIS' Future Moves." *Business Insider*, 2015. <http://www.businessinsider.com/machine-learning-used-to-predict-and-model-isis-2015-9>.
- Lucci, Stephen, and Danny Kopec. *Artificial Intelligence in the 21st Century: A Living Introduction*, 2016.
- Luger, George F., and William A. Stubblefield. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*. 2nd ed. California: Benjamin/Cummings, 1993.
- . *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*. 6th ed., 2008.
- Mack, Chris A. "Fifty Years of Moore's Law." *IEEE Transactions on Semiconductor Manufacturing*

24, no. 2 (2011): 202–7.

Mackenzie. “The Future Military-Artificial Intelligence Complex? | FT Alphaville.” *Financial Times - Alphaville*, 2015. <http://ftalphaville.ft.com/2015/12/15/2147846/the-future-military-artificial-intelligence-complex/>.

Maher, T.M., and S.D. Baum. “Adaptation to and Recovery from Global Catastrophe.” *Sustainability* 5, no. 4 (2013): 1461–79.

Mallory, J. P., and Douglas Q. Adams. *The Oxford Introduction to Proto Indo European and the Proto Indo European World*. New York: Oxford University Press, 2006.

Mamiit, Aaron. “China’s Baidu Unveils New All-Electric, Self-Driving Car: Testing Begins For Modified Chery EQ.” *Tech Times*, August 28, 2016. <http://www.techtimes.com/articles/175317/20160828/chinas-baidu-unveils-new-all-electric-self-driving-car-testing-begins-for-modified-chery-eq.htm>.

Marcus, Gary. “What Comes After the Turing Test?” *The New Yorker*, June 9, 2014. <http://www.newyorker.com/tech/elements/what-comes-after-the-turing-test>.

Markoff, John. “Behind Artificial Intelligence, a Squadron of Bright Real People.” *The New York Times*, October 14, 2005. <http://www.nytimes.com/2005/10/14/technology/behind-artificial-intelligence-a-squadron-of-bright-real-people.html>.

———. “Pentagon Turns to Silicon Valley for Edge in Artificial Intelligence.” *The New York Times*, May 11, 2016. <http://www.nytimes.com/2016/05/12/technology/artificial-intelligence-as-the-pentagons-latest-weapon.html>.

Markoff, John, and Matthew Rosenberg. “China’s Intelligent Weaponry Gets Smarter.” *The New York Times*, February 3, 2017. <https://www.nytimes.com/2017/02/03/technology/artificial-intelligence-china-united-states.html>.

Markov, Igor L. “Limits on Fundamental Limits to Computation.” *Nature* 512, no. 7513 (August 14, 2014): 147–54. doi:10.1038/nature13570.

Marshall, Alex. “From Jingles to Pop Hits, A.I. Is Music to Some Ears.” *The New York Times*, January 22, 2017. <https://www.nytimes.com/2017/01/22/arts/music/jukedeck-artificial-intelligence-songwriting.html>.


Martinage, Robert. “Toward a New Offset Strategy: Exploiting U.s. Long-Term Advantages to Restore U.s. Global Power Projection Capability.” Center for Strategic and Budgetary Assessments, 2014. <http://csbaonline.org/uploads/documents/Offset-Strategy-Web.pdf>.

Mascarenhas, Hyacinth. “Associated Press to Expand Its Sports Coverage by Using AI to Write Minor League Baseball Articles.” *International Business Times UK*, July 5, 2016. <http://www.ibtimes.co.uk/associated-press-expand-its-sports-coverage-by-using-ai-write-minor-league-baseball-articles-1568804>.

McCarthy, J., M.L. Minsky, N. Rochester, and C.E. Shannon. “A Proposal for the Dartmouth

- 
- Summer Research Project on Artificial Intelligence,” August 31, 1955. <http://robotics.cs.tamu.edu/dshell/cs625/2014-09-02.pdf>.
- McCarthy, John, and Stanford University Formal Reasoning Group. “What Is Artificial Intelligence | Basic Questions.” *Formal Reasoning Group*, 2007. <http://www-formal.stanford.edu/jmc/whatisai/node1.html>.
- McCorduck, Pamela. *Machines Who Think*. 2nd ed. Natick, MA: A.K. Peters, Ltd., 2004.
- McCormick, James. “Predictions 2017: Artificial Intelligence Will Drive The Insights Revolution.” Forrester, 2016. https://go.forrester.com/wp-content/uploads/Forrester_Predictions_2017_-_Artificial_Intelligence_Will_Drive_The_Insights_Revolution.pdf.
- McCulloch, Warren S., and Walter Pitts. “A Logical Calculus of the Ideas Immanent in Nervous Activity.” *Bulletin of Mathematical Biophysics* 5 (1943): 115–33.
- McDonald, Coby. “A.I. Downs Expert Human Fighter Pilot In Dogfight Simulation | Popular Science.” *Popular Science*, June 27, 2016. <http://www.popsci.com/ai-pilot-beats-air-combat-expert-in-dogfight?src=SOC&dom=tw>.
- McKinsey Global Institute, Michael Chui, James Manyika, Jacques Bughin, Richard Dobbs, Charles Roxburgh, Hugo Sarrazin, Geoffrey Sands, and Magdalena Westergren. *The Social Economy: Unlocking Value and Productivity Through Social Technologies*. McKinsey Global Institute, 2012.
- McMahan, H. Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. “Communication-Efficient Learning of Deep Networks from Decentralized Data.” *arXiv:1602.05629 [Cs]*, February 17, 2016. <http://arxiv.org/abs/1602.05629>.
- McNeill, William H. *The Pursuit of Power: Technology, Armed Force, and Society since A.D. 1000*. 1 edition. Chicago: University Of Chicago Press, 1984.
- Medeiros, João. “The Startup Fighting Cancer with AI.” *WIRED UK*. Accessed February 25, 2017. <http://www.wired.co.uk/article/ai-cancer-drugs-berg-pharma-startup>.
- Merchant, Brian. “The Poem That Passed the Turing Test.” *Motherboard*, 2015. https://motherboard.vice.com/en_us/article/the-poem-that-passed-the-turing-test.
- Metz, Cade. “AI Is About to Learn More Like Humans—with a Little Uncertainty.” *WIRED*, 2017. <https://www.wired.com/2017/02/ai-learn-like-humans-little-uncertainty/>.
- . “AI’s Factions Get Feisty. But Really, They’re All on the Same Team.” *WIRED*, 2017. <https://www.wired.com/2017/02/ais-factions-get-feisty-really-theyre-team/>.
- . “Google Says Its AI Catches 99.9 Percent of Gmail Spam.” *WIRED*, July 9, 2015. <http://www.wired.com/2015/07/google-says-ai-catches-99-9-percent-gmail-spam/>.
- . “Google’s Hand-Fed AI Now Gives Answers, Not Just Search Results.” *WIRED*, 2016. <https://www.wired.com/2016/11/googles-search-engine-can-now-answer-questions-human-help/>.

- . “Inside the Poker AI That Out-Bluffed the Best Humans.” *WIRED*, 2017. <https://www.wired.com/2017/02/libratus/>.
- . “Teaching Computers to Play Atari Is A Big Step Toward Bringing Robots Into the Real World.” *WIRED*, 2015. <https://www.wired.com/2015/12/teaching-ai-to-play-atari-will-help-robots-make-sense-of-our-world/>.
- Metz, Cage. “Google’s Training Its AI to Be Android’s Security Guard | WIRED.” *WIRED*, May 2, 2016. <http://www.wired.com/2016/06/googles-android-security-team-turns-machine-learning/>.
- Metz, Rachel. “AI Startup Vicarious Claims Victory Over Captchas.” *MIT Technology Review*, 2013. <https://www.technologyreview.com/s/520581/ai-startup-says-it-has-defeated-captchas/>.
- . “Toyota Investing \$50M for Autonomous-Car Research at Stanford and MIT.” *MIT Technology Review*, 2015. <https://www.technologyreview.com/s/541046/toyota-investing-50m-with-stanford-mit-for-autonomous-car-research/>.
- Ministerie van defensie [Ministry of Defence, The Netherlands]. “Houvast in een onzekere wereld Lijnen van ontwikkeling in het meerjarig perspectief voor een duurzaam gereede en snel inzetbare krijgsmacht [A Grip in an Uncertain World. Lines of Development for a Sustainably Ready and Quickly Deployable Armed Force].” Rapport. The Hague: Ministry of Defence, February 14, 2017. <https://www.rijksoverheid.nl/documenten/rapporten/2017/02/14/meerjarig-perspectief-krijgsmacht-houvast-in-een-onzekere-wereld>.
- Ministry of Defence, The Netherlands, The Hague Centre for Strategic Studies, and Future Force Conference 2017. “From Partnerships to Ecosystems: Combining Our Efforts for a More Secure World.” *Future Force Conference*. Accessed January 29, 2017. http://futureforceconference.com/wp-content/uploads/2017/04/Web_FFC-rapportage.pdf.
- Ministry of Defence, The Netherlands, and Jeanine A. Hennis-Plasschaert. “Strategische Kennis- En Innovatieagenda 2016-2020. Vóórblijven in Een Onzekere Wereld [Strategic Knowledge and Innovation Agenda 2016-2020. Staying Ahead in an Uncertain World].” The Hague: Ministry of Defence, October 2016
- Ministry of National Defense of the People’s Republic of China. “The Dawn of the Intelligent Military Revolution.” *People’s Liberation Army Daily*, January 28, 2016. http://www.mod.gov.cn/wqzb/2016-01/28/content_4637961.htm.
- Minsky, Marvin, and Seymour Papert. *Perceptrons*. Oxford, England: M.I.T. Press, 1969.
- MIT Tech Review, Emerging Technology from the. “Google DeepMind Teaches Artificial Intelligence Machines to Read.” *MIT Technology Review*, 2015. <https://www.technologyreview.com/s/538616/google-deepmind-teaches-artificial-intelligence-machines-to-read/>.
- . “The Face Detection Algorithm Set to Revolutionize Image Search.” *MIT Technology*

- 
- Review, 2015. <https://www.technologyreview.com/s/535201/the-face-detection-algorithm-set-to-revolutionize-image-search/>.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Daan Riedmiller. "Playing Atari with Deep Reinforcement Learning," 2013.
- Mokyr, Joel. *The Gifts of Athena: Historical Origins of the Knowledge Economy*. Princeton, [N.J.]: Princeton University Press, 2002.
- Moravec, Hans. *Mind Children: The Future of Robot and Human Intelligence*. 4. print. Cambridge: Harvard Univ. Press, 1995.
- Morris, Ian. *War! What Is It Good for? Conflict and the Progress of Civilization from Primates to Robots*. New York: Picador USA, 2015.
- . *Why The West Rules - For Now: The Patterns of History and What They Reveal about the Future*. Profile Books, 2010.
- Moz. "Google Algorithm Change History - Moz." *Moz*. Accessed September 2, 2016. <https://moz.com/google-algorithm-change>.
- Muehlhauser, Luke. "When Will AI Be Created?" *Machine Intelligence Research Institute*, May 15, 2013. <https://intelligence.org/2013/05/15/when-will-ai-be-created/>.
- . "When Will AI Be Created?" *Machine Intelligence Research Institute*, May 15, 2013. <https://intelligence.org/2013/05/15/when-will-ai-be-created/>.
- Müller, Vincent C., and Nick Bostrom. "Future Progress in Artificial Intelligence: A Survey of Expert Opinion." In *Fundamental Issues of Artificial Intelligence*, 553–570. Springer, 2016. http://link.springer.com/chapter/10.1007/978-3-319-26485-1_33.
- Musthaler, Linda. "How to Use Deep Learning AI to Detect and Prevent Malware and APTs in Real-Time." *Network World*, March 11, 2016. <http://www.networkworld.com/article/3043202/security/how-to-use-deep-learning-ai-to-detect-and-prevent-malware-and-apt-in-real-time.html>.
- NarrativeScience. "Outlook on Artificial Intelligence in the Enterprise - 2016." NarrativeScience, 2016. <http://resources.narrativescience.com/h/i/272429469-outlook-on-artificial-intelligence-in-the-enterprise-2016-research-report>.
- NATO Allied Command Transformation. *Visions of Warfare 2036*. Edited by Trina Marie Phillips and August Cole. Norfolk VA: NATO Allied Command Transformation, 2017.
- Nazre, Ajit, and Rahul Garg. "A Deep Dive in the Venture Landscape of Artificial Intelligence and Machine Learning." September 2015. <http://slideplayer.com/slide/7002258/>.
- Nelson, Richard R. "Physical and Social Technologies and Their Evolution." *Économie Appliquée* 56, no. 3 (2003): 13–32.
- netmarketshare.com. "Search Engine Market Share." *Netmarketshare*, July 2016. <https://www.netmarketshare.com/search-engine-market-share.aspx?qprid=4&qpcustomd=0>.

- Newman, Daniel. "Inside Look: The World's Largest Tech Companies Are Making Massive AI Investments." *Forbes*, 2017. <http://www.forbes.com/sites/danielnewman/2017/01/17/inside-look-the-worlds-largest-tech-companies-are-making-massive-ai-investments/>.
- Newman, LilHay. "AI Can Recognize Your Face Even If You're Pixelated." *WIRED*, 2016. <https://www.wired.com/2016/09/machine-learning-can-identify-pixelated-faces-researchers-show/>.
- Nilsson, Nils J. *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge ; New York: Cambridge University Press, 2010.
- Niu, Jiqiang, Wenwu Tang, Feng Xu, Xiaoyan Zhou, and Yanan Song. "Global Research on Artificial Intelligence from 1990–2014: Spatially-Explicit Bibliometric Analysis." *ISPRS International Journal of Geo-Information* 5, no. 5 (May 16, 2016): 66. doi:10.3390/ijgi5050066.
- Norberg, Johan. *Progress: Ten Reasons to Look Forward to the Future*. Oneworld Publications, 2017.
- nortwolf_sam. "Роботизированный Комплекс «Платформа-М» На Вооружении Тихоокеанского Флота." *Nortwolf_sam*, October 11, 2015. <http://nortwolf-sam.livejournal.com/1054499.html>.
- Nunes, Eric, Ahmad Diab, Andrew Gunn, Ericsson Marin, Vineet Mishra, Vivin Paliath, John Robertson, Jana Shakarian, Amanda Thart, and Paulo Shakarian. "Darknet and Deepnet Mining for Proactive Cybersecurity Threat Intelligence." *arXiv Preprint arXiv:1607.08583*, 2016. <http://arxiv.org/abs/1607.08583>.
- Office of Science and Technology Policy. "The National Artificial Intelligence Research and Development Strategic Plan." National Science and Technology Council, 2016. https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf.
- Oxford Internet Institute. *Video: Politics, Propaganda, and Bots—The Changing Nature of Cyber Warfare*, 2016. <https://www.oii.ox.ac.uk/video-politics-propaganda-and-bots-the-changing-nature-of-cyber-warfare/>.
- Palmer, Jason. "Here, There and Everywhere. Quantum Technology Is Beginning to Come into Its Own." *Economist*, March 9, 2017. <http://www.economist.com/technology-quarterly/2017-03-09/quantum-devices>.
- Parker, Geoffrey, Marshall Van Alstyne, and Sangeet Paul Choudary. *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*, 2016. <https://www.overdrive.com/search?q=747D962F-32B3-4F05-9126-A09E659AB96C>.
- Pavluk, Joshua, and August Cole. "From Strategy to Execution: Accelerating the Third Offset." *War on the Rocks*, June 9, 2016. <http://warontherocks.com/2016/06/from-strategy-to>

execution-accelerating-the-third-offset/.

- Pearl, Judea. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.
- People's Republic of China. "13th Five-Year National Science and Technology Innovation Plan," 2016. http://www.gov.cn/zhengce/content/2016-08/08/content_5098072.htm.
- . "The People's Republic of China National Economic and Social Development of the Thirteenth Five-Year Plan," 2016. http://www.gov.cn/xinwen/2016-03/17/content_5054992.htm.
- Perez, Carlota. *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*. Cheltenham UK; Northampton MA USA: Edward Elgar Publishing, 2002.
- Petrov, Slav. "Announcing SyntaxNet: The World's Most Accurate Parser Goes Open Source." *Google Research Blog*, May 12, 2016. <https://research.googleblog.com/2016/05/announcing-syntaxnet-worlds-most.html>.
- Pinker, Steven. *The Better Angels of Our Nature: Why Violence Has Declined*. Penguin, 2011.
- Plebe, Alessio, and Pietro Perconti. "The Slowdown Hypothesis." In *Singularity Hypotheses*, edited by Amnon H. Eden, James H. Moor, Johnny H. Søraker, and Eric Steinhart, 349–65. The Frontiers Collection. Springer Berlin Heidelberg, 2012. doi:10.1007/978-3-642-32560-1_17.
- Poole, David, Alan Mackworth, and Randy Goebel. *Computational Intelligence: A Logical Approach*. New York: Oxford University Press, 1998. <http://www.cs.ubc.ca/~poole/ci.html>.
- Posner, Richard A. *Catastrophe: Risk and Response*. Oxford; New York: Oxford University Press, 2005.
- Powell, Joanne L., Penelope A. Lewis, Robin I. M. Dunbar, Marta García-Fiñana, and Neil Roberts. "Orbital Prefrontal Cortex Volume Correlates with Social Cognitive Competence." *Neuropsychologia* 48, no. 12 (October 2010): 3554–62. doi:10.1016/j.neuropsychologia.2010.08.004.
- Prahalad, C. K., and Gary Hamel. *The Core Competence of the Corporation*. Harvard Business Review, 2001.
- Pruetz, Jill D., and Paco Bertolani. "Savanna Chimpanzees, Pan Troglodytes Verus, Hunt with Tools." *Current Biology* 17, no. 5 (March 2007): 412–17. doi:10.1016/j.cub.2006.12.042.
- Purdy, Mark, and Paul Daugherty. "Why Artificial Intelligence Is the Future of Growth." Accenture, 2017. https://www.accenture.com/t00010101T000000_w_/gb-en/_acnmedia/PDF-33/Accenture-Why-AI-is-the-Future-of-Growth.PDF#zoom=50.
- Pyne, Stephen J., Greg Bankoff, Uwe Lübken, and Jordan Sand. *Flammable Cities: Urban Conflagration and the Making of the Modern World*. University of Wisconsin Pres, 2012.

- <https://books.google.nl/>
- Reed, Harry L., Jr. "Firing Table Computations on the Eniac." In *Proceedings of the 1952 ACM National Meeting (Pittsburgh)*, 103–106. ACM '52. New York, NY, USA: ACM, 1952. doi:10.1145/609784.609796.
- Reilly, M.B. "Beyond Video Games: New Artificial Intelligence Beats Tactical Experts in Combat Simulation." *University of Cincinnati Magazine*, June 27, 2016. http://magazine.uc.edu/editors_picks/recent_features/alpha.
- Rich, E., and K. Knight. *Artificial Intelligence*. 2nd ed. New York: McGraw-Hill, 1991.
- Rieffel, Eleanor G., and Wolfgang H. Polak. *Quantum Computing: A Gentle Introduction*. 1 edition. Cambridge, Mass: The MIT Press, 2011.
- Rohling, Gitta. "Artificial Intelligence: Facts and Forecasts - Boom for Learning Systems." Siemens, 2014. <http://www.siemens.com/innovation/en/home/pictures-of-the-future/digitalization-and-software/artificial-intelligence-facts-and-forecasts.html>.
- Rosenberg, Karen R., and Wenda R. Trevathan. "Are Human Infants Altricial?," 2015. <http://meeting.physanth.org/program/2015/session15/rosenberg-2015-are-human-infants-altricial.html>.
- Rumelhart, David E., James L. McClelland, and CORPORATE PDP Research Group, eds. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1: Foundations*. Cambridge, MA, USA: MIT Press, 1986.
- Russell, S.J., and Peter Norvig. *Artificial Intelligence: A Modern Approach*, 2009. <http://www.cin.ufpe.br/~tfl2/artificial-intelligence-modern-approach.9780131038059.25368.pdf>.
- Russell, Stuart. "Defining Intelligence." *EDGE*, February 7, 2017. https://www.edge.org/conversation/stuart_russell-defining-intelligence.
- Russell, Stuart, Daniel Dewey, and Max Tegmark. "Research Priorities for Robust and Beneficial Artificial Intelligence." *AI Magazine*, 2015. https://futureoflife.org/data/documents/research_priorities.pdf?x90991.
- Russia Today. "Russia's Mine-Clearing Uran-6 Robots to Help Get Rid of Hidden Explosives in Palmyra (VIDEOS)." *RT International*, March 30, 2016. <https://www.rt.com/news/337810-russia-palmyra-demine-robot/>.
- "Russian Pundit Interviewed on Latest Combat Equipment." *BBC Monitoring Former Soviet Union*, August 4, 2016. <http://global.factiva.com/redirect/default.aspx?P=sa&an=BBCSUP0020160804ec84000m9&cat=a&ep=ASE>.
- Sagan, Carl. "Nuclear War and Climatic Catastrophe: Some Policy Implications." *Foreign Affairs* 62, no. 2 (1983): 257–92. doi:10.2307/20041818.
- Samuel, Arthur L. "Some Studies in Machine Learning Using the Game of Checkers." *IBM Journal* 3, no. 3 (July 1959). <https://www.cs.virginia.edu/~evans/greatworks/samuel1959>.




pdf.


- Schalkoff, R.I. *Artificial Intelligence: An Engineering Approach*. New York: McGraw-Hill, 1990.
- Scharre, Paul. "Robotics on the Battlefield, Part I: Range, Persistence and Daring." 20YY Series: Preparing for War in the Robotic Age. Center for a New American Security, May 2014. http://www.cnas.org/sites/default/files/publications-pdf/CNAS_RoboticsOnTheBattlefield_Scharre.pdf.
- Schatzberg, Eric. "Technik Comes to America: Changing Meanings of Technology before 1930." *Technology and Culture* 47, no. 3 (2006): 486–512.
- Schmidt, Eric. *Google: How Google Works*. First edition. New York: Grand Central Publishing, 2014.
- Scholl, Zackary. "Turing Test: Passed, Using Computer-Generated Poetry." *Raspberry Pi AI*, January 24, 2015. <https://rpi.ai.wordpress.com/2015/01/24/turing-test-passed-using-computer-generated-poetry/>.
- Schwab, Klaus. "The Fourth Industrial Revolution." World Economic Forum Geneva, 2016. <https://foreignaffairs.org/articles/2015-12-12/fourth-industrial-revolution>.
- . *The Fourth Industrial Revolution*. World Economic Forum, 2016.
- Senor, Dan, and Saul Singer. *Start-up Nation: The Story of Israel's Economic Miracle*. Reprint edition. New York: Twelve, 2011.
- Seyfarth, Robert M., and Dorothy L. Cheney. "Affiliation, Empathy, and the Origins of Theory of Mind." *PNAS* 110, no. 2 (2013). http://m.pnas.org/content/110/Supplement_2/10349.full.pdf.
- Shalal, Andrea. "Pentagon Eyes \$12-15 Billion for Early Work on New Technologies." *Reuters*, December 14, 2015. <http://www.reuters.com/article/us-usa-military-technologies-idUSKBN0TX1UR20151214>.
- Shokri, Reza, and Vitaly Shmatikov. "Privacy-Preserving Deep Learning," 1310–21. ACM Press, 2015. doi:10.1145/2810103.2813687.
- Shorey, Samantha, and Philip N. Howard. "Automation, Algorithms, and Politics| Automation, Big Data and Politics: A Research Review." *International Journal of Communication* 10, no. 0 (October 12, 2016): 24.
- Shumaker, Robert W., Kristina R. Walkup, and Benjamin B. Beck. *Animal Tool Behavior: The Use and Manufacture of Tools by Animals*. JHU Press, 2011.
- Simonite, Tom. "A Cloud That Can't Leak." *MIT Technology Review*, 2011. <https://www.technologyreview.com/s/424942/a-cloud-that-cant-leak/>.
- . "Microsoft and Google Want to Let Artificial Intelligence Loose on Our Most Private Data." *MIT Technology Review*, 2016. <https://www.technologyreview.com/s/601294/microsoft-and-google-want-to-let-artificial-intelligence-loose-on-our-most-private->




data/.

- Singh, Abhijit. "Is China Really Building Missiles With Artificial Intelligence?" *The Diplomat*, 2016. <http://thediplomat.com/2016/09/is-china-really-building-missiles-with-artificial-intelligence/>.
- Sobradillo, Patricia, Francisco Pozo, and Álar Agustí. "P4 Medicine: The Future around the Corner." *Archivos de Bronconeumología ((English Edition))* 47, no. 1 (2011): 35–40.
- Sputnik. "Robots on Patrol: Russian Borders to Be Guarded by Artificial Intelligence," 2016. <https://sputniknews.com/military/201606301042223143-russian-borders-guarded-robots/>.
- Stanford University. "Artificial Intelligence and Life in 2030 One Hundred Year Study on Artificial Intelligence | Report of the 2015 Study Pane." *Stanford One Hundred Year Study on Artificial Intelligence (AI100)*, September 2016. https://ai100.stanford.edu/sites/default/files/ai_100_report_0901fnlc_single.pdf.
- Stanton, Andrew, Amanda Thart, Ashish Jain, Priyank Vyas, Arpan Chatterjee, and Paulo Shakarian. "Mining for Causal Relationships: A Data-Driven Study of the Islamic State." *arXiv:1508.01192 [Cs]*, August 5, 2015. <http://arxiv.org/abs/1508.01192>.
- statista.com. "Mobile Phone Users Worldwide 2013-2019." *Statista*, 2016. <https://www.statista.com/statistics/274774/forecast-of-mobile-phone-users-worldwide/>.
- Sullivan, Danny. "FAQ: All about the Google RankBrain Algorithm." *Search Engine Land*, June 23, 2016. <http://searchengineland.com/faq-all-about-the-new-google-rankbrain-algorithm-234440>.
- Sutherland, Stephani. "What Happens in the Brain When We Read?" *Scientific American Mind* 26, no. 4 (August 7, 2015): 14–14.
- TensorFlow. "TensorFlow — an Open Source Software Library for Machine Intelligence." *TensorFlow*, 2016. <https://www.tensorflow.org/>.
- Tetlock, Philip E., and Dan Gardner. *Superforecasting: The Art and Science of Prediction*. First edition. New York: Crown Publishers, 2015.
- The Economist. "Anti-Submarine Warfare: Seek, but Shall Ye Find?" *The Economist*, August 6, 2016. <http://www.economist.com/news/science-and-technology/21703360-proliferation-quieter-submarines-pushing-navies-concoct-better-ways>.
- . "Are You Talking to Me? -." *The Economist*, June 7, 2007. <http://www.economist.com/node/9249338>.
- . "Machine Learning: Of Prediction and Policy," August 20, 2016. <http://www.economist.com/news/finance-and-economics/21705329-governments-have-much-gain-applying-algorithms-public-policy?frsc=dgla>.
- . "The Future of War: Israel First to Deploy Fully Automated Military Robots."

- 
- The Mainichi*, August 24, 2016. <http://global.factiva.com/redirect/default.aspx?P=sa&an=AIWMDM0020160824ec800105&cat=a&ep=ASE>.
- The Guardian. "Computer Simulating 13-Year-Old Boy Becomes First to Pass Turing Test." *The Guardian*, June 9, 2014, sec. Technology. <https://www.theguardian.com/technology/2014/jun/08/super-computer-simulates-13-year-old-boy-passes-turing-test>.
- The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*, 2014.
- Timberg, Craig. "Russian Propaganda Effort Helped Spread 'fake News' during Election, Experts Say." *Washington Post*, 2016. https://www.washingtonpost.com/business/economy/russian-propaganda-effort-helped-spread-fake-news-during-election-experts-say/2016/11/24/793903b6-8a40-4ca9-b712-716af66098fe_story.html.
- Toffler, Alvin. *The Third Wave*. A Bantam Book. New York: Bantam Books, 1990.
- Toffler, Alvin, and Heidi Toffler. *War and Anti-War Survival at the Dawn of the 21st Century*. New York: Little Brown, 1993.
- Topol, Sarah A. "Killer Robots Are Coming And These People Are Trying To Stop Them." *BuzzFeed*. Accessed March 21, 2017. <https://www.buzzfeed.com/sarahatopol/how-to-save-mankind-from-the-new-breed-of-killer-robots>.
- Trabulsi, Andrew. "Future of Artificial Intelligence." presented at the Institute for the Future, June 2015. http://www.fujitsu.com/us/Images/Panel1_Andrew_Trabulsi.pdf.
- Tractica. "Early Enterprise Adopters of Artificial Intelligence Will Include the Business Services, Advertising, Media & Entertainment, and Finance & Investment Sectors." *Tractica*, 2016. <https://www.tractica.com/newsroom/press-releases/early-enterprise-adopters-of-artificial-intelligence-will-include-the-business-services-advertising-media-entertainment-and-finance-investment-sectors/>.
- Trappenberg, Thomas. *Fundamentals of Computational Neuroscience*. 2 edition. Oxford ; New York: Oxford University Press, 2010.
- Turing, A.M. "Computing Machinery and Intelligence." *Mind* 59, no. 236 (October 1950): 433–60.
- United States, Joint Chiefs of Staff. *Foreign Internal Defense*. Joint Publication, JP 3-22. Washington, D.C.: Joint Chiefs of Staff, 2010. <http://purl.fdlp.gov/GPO/gpo29282>.
- United States Marine Corps. *Science Fiction Futures; Marine Corps Security Environment Forecast: Futures 2030-2045*. Quantico, VA: United States Marine Corps, 2016. http://www.mcw.marines.mil/Portals/34/Documents/FuturesAssessment/Marine%20Corps%20Science%20Fiction%20Futures%202016_12_9.pdf?ver=2016-12-09-105855-733.
- University of Cambridge. "Artificially-Intelligent Robot Scientist 'Eve' Could Boost Search for New Drugs." *University of Cambridge*, February 4, 2015. <http://www.cam.ac.uk/research/news/artificially-intelligent-robot-scientist-eve-could-boost-search-for-new-drugs>.

- 
- Urban, Tim. "The Artificial Intelligence Revolution: Part 1." *Wait But Why*, January 22, 2015. <http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html>.
- US Congress. Defense Authorization Act, Pub. L. No. 91-121 (1969).
- US Senate Subcommittee on Space, Science and Competitiveness. The Dawn of Artificial Intelligence [2016]. <http://www.commerce.senate.gov/public/index.cfm/2016/11/commerce-announces-first-artificial-intelligence-hearing>.
- Vardi, Moshe Y. "Artificial Intelligence: Past and Future." *Communications of the ACM* 55, no. 1 (January 1, 2012): 5-5. doi:10.1145/2063176.2063177.
- Veeramachaneni, Kalyan, and Ignacio Araldo. "AI2: Training a Big Data Machine to Defend." Accessed August 20, 2016. http://people.csail.mit.edu/kalyan/AI2_Paper.pdf.
- Venture Scanner. "Artificial Intelligence Q1 Update in 15 Visuals," March 22, 2016. <https://insights.venturescanner.com/2016/03/22/artificial-intelligence-q1-update-in-15-visuals/>.
- Vidra, Eze. "30 Machine Intelligence Startups to Watch in Israel." *Medium*, February 15, 2017. <https://medium.com/@ediggs/30-machine-intelligence-startups-to-watch-in-israel-a05b6597c4a5>.
- Vise, David A., and Mark Malseed. *The Google Story*. New York: Delacorte Press, 2005.
- Vuurst, Louis van de. "Johan Cruijff: Van Trainer Tot Adviseur." *Ajax.nl*, Maart 2016. <http://www.ajax.nl/streams/ajax-actueel/johan-cruijff-van-trainer-tot-adviseur-.htm>.
- Walton, Timothy A. "Securing The Third Offset Strategy: Priorities For Next US Secretary Of Defense – Analysis | CSBA." *CSBA | Center for Strategic and Budgetary Assessments*, July 27, 2016. <http://csbaonline.org/2016/07/27/securing-the-third-offset-strategy-priorities-for-next-us-secretary-of-defense-analysis/>.
- Werbos, Paul. "Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences." Ph.D., Harvard University, 1974.
- Williams, Kevin, Elizabeth Bilsland, Andrew Sparkes, Wayne Aubrey, Michael Young, Larisa N. Soldatova, Kurt De Grave, et al. "Cheaper Faster Drug Development Validated by the Repositioning of Drugs against Neglected Tropical Diseases." *Journal of The Royal Society Interface* 12, no. 104 (March 6, 2015): 20141289. doi:10.1098/rsif.2014.1289.
- Winston, Patrick Henry. *Artificial Intelligence*. 3rd ed., 1992.
- Wong, Kate. "Why Humans Give Birth to Helpless Babies." *Scientific American Blog Network*, August 28, 2012. <https://blogs.scientificamerican.com/observations/why-humans-give-birth-to-helpless-babies/>.
- Woolley, Samuel C., and Philip N. Howard. "Automation, Algorithms, and Politics| Political Communication, Computational Propaganda, and Autonomous Agents — Introduction." *International Journal of Communication* 10, no. 0 (October 12, 2016): 9.
- Xiao, Eva. "The Chinese Government Wants A 100 Billion RMB AI Market By 2018." *TechNode*,

- 
- May 27, 2016. <http://technode.com/2016/05/27/chinese-goverment-wants-100-billion-level-artificial-intelligence-market-2018/>.
- Xinhua News Agency. "中国'脑计划'纳入规划全面展开 坚持'一体两翼'." *Xinhua News Agency*, 2016. http://news.xinhuanet.com/politics/2016-08/18/c_129238381.htm.
- Xiong, Wayne, Jasha Droppo, Xuedong Huang, Frank Seide, Mike Seltzer, Andreas Stolcke, Dong Yu, and Geoffrey Zweig. "Achieving Human Parity in Conversational Speech Recognition." *arXiv Preprint arXiv:1610.05256*, 2016. <https://arxiv.org/abs/1610.05256>.
- Yudelson, Michael V., Kenneth R. Koedinger, and Geoffrey J. Gordon. "Individualized Bayesian Knowledge Tracing Models." *Artificial Intelligence in Education*, 2013, 171–80.
- Yudkowsky, Eliezer. "Artificial Intelligence as a Positive and Negative Factor in Global Risk." *Global Catastrophic Risks* 1 (2008): 303.
- Zhang, Sarah. "China's Artificial-Intelligence Boom." *The Atlantic*, February 16, 2017. <https://www.theatlantic.com/technology/archive/2017/02/china-artificial-intelligence/516615/>.
- Zilis, Shivon. "The Current State of Machine Intelligence 3.0." *Shivon Zilis*, 2016. <http://www.shivonzilis.com//machineintelligence>.
- Грогман, Ян. "Российский Робот-Танк «Уран-9»: Бесплезная Игрушка Или Революция?" *ИноСМИ.Ru*, April 13, 2016. <http://inosmi.ru/military/20160413/236108861.html>.
- Тимошенко, Михаил "Война под микроскопом," *KP.RU - Сайт «Комсомольской Правды»*, August 6, 2015, <http://www.kp.ru/daily/26415/3289072/>
- Птичкин, Сергей. "Умная пуля видит цель." *Российская Газета*, July 19, 2016, sec. СОБЫТИЯ И КОММЕНТАРИИ; Арсенал.
- Толер, Арик. "Использовала Ли Россия Боевых Роботов В Сирии?" *ИноСМИ.Ru*, January 17, 2016. <http://inosmi.ru/military/20160117/235078297.html>.
- ФБА «Экономика сегодня». "Назад В Будущее: Путин Объявил Эру Боевых Роботов." *Рамблер/Новости*, January 26, 2017. <https://news.rambler.ru/politics/35932393-pazad-v-buduschee-putin-obyavil-eru-boevyh-robotov/>. "Эксперт Не Исключил, Что Новейший комплекс 'Уран-9' испытают В Сирии." *РИА Новости*, January 9, 2017. <https://ria.ru/arms/20170109/1485307092.html>.



APPENDIX I: ASILOMAR AI PRINCIPLES



Appendix I: Asilomar AI Principles³²⁶

Research Issues

Research Goal: The goal of AI research should be to create not undirected intelligence, but beneficial intelligence.

Research Funding: Investments in AI should be accompanied by funding for research on ensuring its beneficial use, including thorny questions in computer science, economics, law, ethics, and social studies, such as:

- How can we make future AI systems highly robust, so that they do what we want without malfunctioning or getting hacked?
- How can we grow our prosperity through automation while maintaining people's resources and purpose?
- How can we update our legal systems to be more fair and efficient, to keep pace with AI, and to manage the risks associated with AI?
- What set of values should AI be aligned with, and what legal and ethical status should it have?

Science-Policy Link: There should be constructive and healthy exchange between AI researchers and policy-makers.

Research Culture: A culture of cooperation, trust, and transparency should be fostered among researchers and developers of AI.

Race Avoidance: Teams developing AI systems should actively cooperate to avoid corner-cutting on safety standards.

326 Future of Life Institute, "Asilomar AI Principles."



Ethics and Values

Safety: AI systems should be safe and secure throughout their operational lifetime, and verifiably so where applicable and feasible.

Failure Transparency: If an AI system causes harm, it should be possible to ascertain why.

Judicial Transparency: Any involvement by an autonomous system in judicial decision-making should provide a satisfactory explanation auditable by a competent human authority.

Responsibility: Designers and builders of advanced AI systems are stakeholders in the moral implications of their use, misuse, and actions, with a responsibility and opportunity to shape those implications.

Value Alignment: Highly autonomous AI systems should be designed so that their goals and behaviors can be assured to align with human values throughout their operation.

Human Values: AI systems should be designed and operated so as to be compatible with ideals of human dignity, rights, freedoms, and cultural diversity.

Personal Privacy: People should have the right to access, manage and control the data they generate, given AI systems' power to analyze and utilize that data.

Liberty and Privacy: The application of AI to personal data must not unreasonably curtail people's real or perceived liberty.

Shared Benefit: AI technologies should benefit and empower as many people as possible.

Shared Prosperity: The economic prosperity created by AI should be shared broadly, to benefit all of humanity.

Human Control: Humans should choose how and whether to delegate decisions to AI systems, to accomplish human-chosen objectives.

Non-subversion: The power conferred by control of highly advanced AI systems should respect and improve, rather than subvert, the social and civic processes on which the health of society depends.

AI Arms Race: An arms race in lethal autonomous weapons should be avoided.



Longer-term Issues

Capability Caution: There being no consensus, we should avoid strong assumptions regarding upper limits on future AI capabilities.

Importance: Advanced AI could represent a profound change in the history of life on Earth, and should be planned for and managed with commensurate care and resources.

Risks: Risks posed by AI systems, especially catastrophic or existential risks, must be subject to planning and mitigation efforts commensurate with their expected impact.

Recursive Self-Improvement: AI systems designed to recursively self-improve or self-replicate in a manner that could lead to rapidly increasing quality or quantity must be subject to strict safety and control measures.

Common Good: Superintelligence should only be developed in the service of widely shared ethical ideals, and for the benefit of all humanity rather than one state or organization.

The Hague Centre for Strategic Studies

Lange Voorhout 16

info@hcss.nl

2514 EE The Hague

HCSS.NL

The Netherlands